

RL-TR-97-184
Final Technical Report
October 1997



DESIGN OF AN ADVANCED DEVELOPMENT MODEL OPTICAL DISK- BASED REDUNDANT ARRAY OF INDEPENDENT DISKS (RAID) HIGH SPEED MASS STORAGE SUBSYSTEM

Rising Edge Technologies, Inc.

Michael E. Lewis, Donald M. Maffei, Michael M. Schutte

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

19980414 170

DTIC QUALITY INSPECTED 4

**Rome Laboratory
Air Force Materiel Command
Rome, New York**

This report has been reviewed by the Rome Laboratory Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RL-TR-97-184 has been reviewed and is approved for publication.

APPROVED:

Fred N. Haritatos

FRED N. HARITATOS

Project Engineer

FOR THE DIRECTOR:

J. Camera

JOSEPH CAMERA

Technical Director

Intelligence & Reconnaissance Directorate

If your address has changed or if you wish to be removed from the Rome Laboratory mailing list, or if the addressee is no longer employed by your organization, please notify RL/IRAP, 32 Hangar Road, Rome, NY 13441-4114. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</p>			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	October 1997	Final Jul 94 - Sep 96	
4. TITLE AND SUBTITLE		5. FUNDING NUMBERS	
DESIGN OF AN ADVANCED DEVELOPMENT MODEL OPTICAL DISK-BASED REDUNDANT ARRAY OF INDEPENDENT DISKS (RAID) HIGH SPEED MASS STORAGE SUBSYSTEM		C - F30602-94-C-0046 PE - 63726F PR - 3192 TA - 00 WA - 07	
6. AUTHOR(S)		8. PERFORMING ORGANIZATION REPORT NUMBER	
Michael E. Lewis, Donald M. Maffei, Michael M. Schutte		N/A	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
Rising Edge Technologies, Inc. 1110 Elden Street, Suite 301 Herndon VA 20170-5527		RL-TR-97-184	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		12a. DISTRIBUTION AVAILABILITY STATEMENT	
Rome Laboratory/IRAP 32 Hangar Road Rome NY 13441-4114		Approved for public release; distribution unlimited	
11. SUPPLEMENTARY NOTES		12b. DISTRIBUTION CODE	
Rome Laboratory Project Engineer: Fred N. Haritatos/IRAP(315) 330-4582			
13. ABSTRACT (Maximum 200 words)		<p>The objective of this effort is to design, fabricate, test, deliver and support an Optical Redundant Array of Independent Risks (O-RAID) subsystem for use at Rome Laboratory to demonstrate mass data storage and retrieval operations. In developing new concepts and systems for mass storage, performance and reliability are important features. The O-RAID storage subsystem is designed to enhance the capabilities offered by traditional optical disk drives and RAID storage systems. The O-RAID connects to a variety of host computer systems to provide a high performance, highly reliable mass storage system. The O-RAID design combines the redundant features prescribed to magnetic disk storage with the infinite near-line capabilities, removability and archivability of optical disk technology. The subsystem design derives as an integrated storage product consisting of multiple optical disk drives designed to appear as a single disk to a host computer.</p>	
14. SUBJECT TERMS		15. NUMBER OF PAGES 224	
Computer Storage, Optical Memory, System Architecture		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Table of Contents

ABSTRACT	VIII
SECTION 1 PROGRAM SUMMARY	1
SECTION 2 INTRODUCTION	5
2.1 Background	5
2.2 Significance of ORAID ADM Development	5
2.3 Development Goals	6
2.3.1 Design Approach	6
2.3.2 Host Control Platform	8
2.3.3 Modes of Operation	9
2.3.4 Performance Goals	9
2.3.5 Interface Specifications	10
2.3.6 Maintenance	10
2.4 Contract Scope	10
SECTION 3 ANALYSIS AND EVALUATION	12
3.1 Analysis of Commercial Optical Disk Drives	12
3.1.1 Review of Optical Drive Manufacturers	12
3.1.2 Analysis of Optical Disk Drive Technical Specifications	14
3.1.3 Discussion of Specific Optical Drive Parameters	15
3.1.4 Testing Goals	17
3.1.5 Test Methodology	17
3.2 Analysis of Standard RAID Levels	19
3.3 Analysis of Commercial RAID Controllers	22
3.3.1 RAID Controllers	22
3.3.2 Review of RAID Controller Vendors	22
3.3.3 Analysis of RAID Controllers	24
3.3.4 Sector Size Compatibility	24
3.3.5 Spindle Synchronization	25
3.3.6 System Requirements and RAID Controller Specifications	26
3.4 General System Requirements	26
3.4.1 ORAID Power Requirements	26
3.4.2 Thermal Considerations	27
SECTION 4 RESULTS	28
4.1 Results of Optical Drive Analysis	28
4.1.2 Test Configuration	30
4.1.3 Optical Drive Testing	31
4.1.4 Write/Read Performance	33
4.1.5 Tabulation and Recording	35
4.1.6 Results	35
4.1.7 Comparative Analysis of Tested Drives	42

4.1.8 Optical Disk Drives: Lessons Learned	44
4.2 Results of RAID Controller Analysis	45
4.2.1 RAID Controller Vendors	45
4.2.2 RAID Controller Testing	46
4.2.3 Controller Evaluation and Selection	49
4.2.4 OEM RAID Controllers: Lessons Learned	51
4.3 System Design Considerations	52
4.3.1 Evaluation of Design Options	53
4.3.2 Mechanical Design Considerations	55
4.3.3 ORAID Power Supplies	57
4.3.4 System Design: Lessons Learned	58
4.4 Selection of an Integration Approach	59
4.4.1 Embedded Controller Evaluation and Selection	59
4.4.2 System Interface Component Design and Development	60
4.4.3 Chassis Integration	63
4.4.5 Lessons Learned	66
SECTION 5 ANALYSIS OF FINAL PRODUCT	67
5.1 Comparison of Results Against Development Goals	67
5.1.1 Optical Disk Storage Subsystem	67
5.1.2 RAID Controller	67
5.1.3 Equipment Chassis	67
5.1.4 Host Control Platform	68
5.1.5 Modes of Operation	68
5.1.6 Performance Goals	69
5.1.7 Interface Specifications	70
5.1.8 Maintenance	70
5.2 Final Analysis	70
SECTION 6 REDESIGN AND REFINEMENT	72
6.1 Changes in the Optical Drive Market	72
6.1.1 Selecting an Optical Drive	72
6.1.2 Testing	72
6.2 Adapting Lessons Learned	73
6.2.1 Optical Disk Drives	73
6.2.2 OEM RAID Controller	74
6.2.3 System Design	74
6.2.4 System Interface	76
6.3 Analysis of Final Product	76
6.4 Further Development and Enhancements	77
APPENDIX: OPTICAL DRIVE TEST RESULTS	79/80

List of Figures

SECTION 3 ANALYSIS AND EVALUATION

Figure 3- 1. Spindle Synchronization	25
--------------------------------------	----

SECTION 4 RESULTS

Figure 4- 1. Calculated Transfer Rate, 1x Drive	29
Figure 4- 2. Calculated Transfer Rate, 2x Drive	29
Figure 4- 3. Calculated Transfer Rate, 3x Drive	30
Figure 4- 4. Basic Disk Test Configuration	31
Figure 4- 5. Average Read Rate, Olympus MOS525E	36
Figure 4- 6. Average Write Rate, Olympus MOS525E	36
Figure 4- 7. Average Read Rate, HPC1716T	37
Figure 4- 8. Average Write Rate, HPC1716T	38
Figure 4- 9. Average Read Rate, Hitachi OD152	39
Figure 4- 10. Average Write Rate, Hitachi OD152	39
Figure 4- 11. Average Read Rate, Maxoptics T3	40
Figure 4- 12. Average Write Rate, Maxoptics T3	41
Figure 4- 13. Average Read Rate, Pinnacle Sierra	42
Figure 4- 14. Average Write Rate, Pinnacle Sierra	42
Figure 4- 15. Preliminary Results of Integration Testing	53
Figure 4- 16. Removable Drive Tray	56
Figure 4- 17. System Block Diagram	61

List of Tables

SECTION 2 INTRODUCTION

Table 2-1. Benefits of Optical Storage and RAID Systems	6
Table 2-2. ORAID Program Development Model	11

SECTION 3 ANALYSIS AND EVALUATION

Table 3-1. Optical Disk Drive Manufacturers	13
Table 3-2. Summary of Optical Drive Requirements	14
Table 3-3. Summary of Selected Optical Drives	15
Table 3-4. Optical Drive Performance Tests	17
Table 3-5. Drive Test Matrix	19
Table 3-6. Summary of RAID Levels	19
Table 3-7. Analysis of RAID Features	21
Table 3-8. RAID Controller Vendors	23
Table 3-9. Controller Product Specifications	24
Table 3-10. Power Supply Vendors	27

SECTION 4 RESULTS

Table 4-1. Selected Optical Drive	28
Table 4-2. Switch Selectable or Software Selectable Commands	32
Table 4-3. Set-Up and Verification Test Commands	33
Table 4-4. SCSI Commands for Performance Testing	34
Table 4-5. Key Factors Affecting Drive Performance	34
Table 4-6. Drive Comparison Criteria.	43
Table 4-7. Weighted Drive Comparison.	44
Table 4-8. Controller Product Specifications	46
Table 4-9. Analysis of Concept 810	47
Table 4-10. Analysis of Z9000	48
Table 4-11. Analysis of CRD-5000	49
Table 4-12. Description of Evaluation Criteria	49
Table 4-13. Controller Comparison Criteria	50
Table 4-14. Areas of Concern for COTS System Integration	50
Table 4-15. Digi-Data Z9000 Revisions	51
Table 4-16. Design Approaches	53
Table 4-17. Power Supply Vendors	57
Table 4-18. Preliminary Power Budget	58
Table 4-19. Evaluation of PC/104 Products	60
Table 4-20. Embedded Code Modules	60

Table 4-21. Developed Components	62
Table 4-22. Fabricated Parts, ORAID ADM Chassis	64

SECTION 5 ANALYSIS OF FINAL PRODUCT

Table 5-1. Performance Goals and Results	69
--	----

SECTION 6 REDESIGN AND REFINEMENT

Table 6-1. Fabricated Chassis Components	74
Table 6-2. List of ORAID ADM Assemblies	75
Table 6-3. Performance Goals and Results	77

Abstract

The Advanced Development Model (ADM) of the Optical disk-based RAID (ORAID) was intended to evaluate the feasibility of integrating commercial optical disk drives and RAID components into a single, high performance storage unit. The design was partitioned into three basic elements; the optical disk drives, the rewritable media, and the electronics controller. Test and evaluation of commercial components was performed in each of these categories in order to characterize performance and determine applicability toward the successful integration of a final unit.

With a thorough understanding of commercially available products in hand, a detailed system design and analysis was performed. The system design utilized analysis and integration to accomplish several goals. Analysis was used to investigate and select a RAID configuration, optical disk drives, and a RAID controller that would be compatible with the system performance requirements. Integration applied physical component testing with the analysis to create a final system design.

As a result of these efforts, the best combination of commercial components along with required custom designed hardware and software created a highly functional system. This combination of system elements resulted in development of a fully integrated ORAID that met the goals of the ORAID ADM.

Section 1 Program Summary

The Advanced Development Model (ADM) of an Optical disk-based Redundant Array of Inexpensive Disks (OR RAID) was intended to determine the feasibility of developing a single, integrated, high performance unit capable of combining the benefits of optical disk technology within a RAID system. The system design proceeded along a path of analysis of commercial components, testing of select commercial components, development of required components, and final system integration and testing.

The analysis process began with a review of commercial product offerings from optical disk drive manufacturers. This study revealed that there are a handful of vendors providing optical disk drives that met the basic requirements of the OR RAID. The basic functions required by the system could be met with any of the drives that were evaluated. However, choosing a specific unit also effects the final OR RAID architecture. For example, using specified drive transfer rate as a selection criteria, the system transfer rate could be satisfied by four optical drives. However when viewed from a capacity perspective, 10 or 12 drives would be required to meet system specifications. This implies a RAID controller for the system that must support 11 or 13 drives (if a parity drive were included). Thus, using a single factor in drive selection could present other problems in system design.

Drive analysis and testing was conducted in order to select a commercially available optical disk drive that would provide the best balance between read/write data transfer performance and capacity. In order to select the best drive in this category, optical drives from several manufacturers were extensively tested. The results of these tests indicated the Maxoptics T3-1300 with 4 MB of cache (TMT3/4) would provide the best performance for the current generation of OR RAID products. The next phase of the analysis was conducted to determine the proper RAID level that would provide the best system performance.

Information was collected and an analysis performed of the salient characteristics of all available RAID levels. Based on this research, a RAID Level 3 architecture was selected to meet the requirements of the OR RAID contract. All of the basic data acquisition and storage parameters can be met and sufficient commercial vendor support exists for products supporting the Level 3 architecture. The next step in the analysis evaluated RAID controllers. Results of the controller analysis were combined with optical drive analysis in order to determine the specifications of the resulting system.

A review of commercial RAID controller manufacturers revealed that there are many RAID system suppliers but few vendors of the actual RAID controller electronics. Several options for OEM-style RAID controllers were discovered that met the basic requirements of the OR RAID. On the surface, any of the controllers evaluated appear to meet the basic functional requirements of the system. However, careful analysis and testing quickly reduced the number of viable candidates. Selection of a specific unit also

determined the final architecture and thus its ability to meet the design goals of the ORAID ADM.

For example, CMD Incorporated offers a controller capable of supporting six data drives. Utilizing this controller yields a single side storage capacity of 3.9 GBytes (using 1.3 GByte drives) or 6 GBytes (using 2 GBytes drives). Higher storage capacities could be obtained using the eight data drive controller utilized in Ciprico, Incorporated or Storage Concepts, Incorporated RAID systems. The selected RAID controller must also be viewed from a transfer rate perspective. As stated earlier, using a published transfer rates yields a six drive system with 6.6 Mbytes/second. This represents a sustained transfer rate that is more than double the required sustained system transfer rate. Therefore, a study of possible controller and drive combinations and an associated trade study was performed.

The selected disk drive was then coupled with possible RAID controllers to create various system architectures. Trade studies and discussions with the controller manufacturers were then conducted. The information was used to create a system architecture that achieved as many of the ORAID ADM specifications as possible. Based on possible system configurations, RAID controllers were tested for use in the system.

The evaluation of available commercial RAID controllers resulted in three being selected for testing. These three units were the only controllers that met the requirements of the ORAID ADM. The Digi-Data, Incorporated Z9000, Storage Concepts, Incorporated Concept 810, and the CMD, Incorporated CRD-5000 controllers were tested. Preliminary testing was performed to establish compatibility of the base controller with the optical drives. The test includes the ability to set up the controller for different media, transfer rates, and spin up times, as well as the ability of the controller to recognize the optical drives on initial power up.

After controller testing was completed, a weighted comparison of controller characteristics was performed. The evaluation criteria were established based on factors that were important to the successful integration of a commercial RAID controller with optical drives. A weight was assigned to each factor with respect to its importance in determining an overall score for each controller. The test results for each controller were then analyzed against each measurement criteria. Based on test results and comparative analysis, the Digi-Data Z9000 was the clear winner. This is the controller that was used to implement the ORAID ADM.

With the drive and RAID controller selected, basic integration testing and evaluation of system compliance with design goals was performed. The results of this initial level of testing pointed out some limitations in the integrated design. Therefore, a review of the integrated approach was conducted to evaluate ways to overcome some of these limitations. This evaluation resulted in creating a hybrid approach consisting of the commercial components and custom designed hardware and software.

Through analysis of current technology as well as development of custom components, a complete system was developed that closely mirrored the program goals.

As the first article ADM was being prepared for delivery, a dramatic change occurred in the optical drive market. All major industry players announced the release of their next generation optical drives. Capacity doubled, the drive was repackaged as a half height device, and transfer rates increased. These changes represent a marked improvement in the capabilities of the ORAID ADM. However, these drive changes also forced several major design changes to the ORAID ADM.

After delivery of the first article ADM, changes were implemented in the basic ADM concept to accommodate the new drives. Because the ORAID ADM was developed around the industry standard SCSI architecture, the necessary changes were largely mechanical and package related.

A cursory review of the optical drive market was performed. Three basic drive types were seen emerging, standard, LIM-DOW, and high capacity. Standard drives offered 2.6 GBytes of storage (1.3 GBytes per side) and transfer rates up to 4 MB/second. One inherent problem with optical storage has traditionally been the two pass write (one pass erases, one pass writes with an optional verify pass). Light Intensity Modulation Direct Overwrite (LIM-DOW) eliminates the two pass write. Thus, standard 2.6 GByte drives can read and write data at the same speed. High capacity drives offer 4.6 GBytes of storage (2.3 GBytes per side).

Standard drives are available from a number of vendors. Nikon Precision Incorporated developed the LIM-DOW technology for their own use and are currently licensing the technology to Most and Plasmon (others plan to license the technology in the near future). The high capacity drives are currently only offered by Pinnacle Micro as their Apex drive and require non-standard media. Both the LIM-DOW technology and Apex were announced well before all of the design and production issues had been solved by their respective manufacturers. Although each of these drives were tested, they were very early versions of the final product and unusable for this second generation ORAID ADM. Future generations of the ORAID will be able to take advantage of these new technologies. A standard 2.6 GByte optical drive from Maxoptics will be used for the second generation ORAID ADM.

The reduction in height of the drives created the largest single change to the basic ORAID ADM. The original packaging was designed to support full height drives with no consideration given to supporting smaller units. Therefore, a significant amount of package redesign and improvement needed to be performed to complete delivery of the second ORAID ADM. Packaging was improved to accommodate the half-height drives and improve the hot-swappability of system components.

The final improvement in the second generation ORAID ADM involved an upgrade to the user interface. The first article ADM relied on a PC/104 CPU card in conjunction with a custom designed circuit card to perform system control, monitoring, and display driving functions. Although functional, there were many limitations to this approach. Therefore, to improve system operation, these components were removed and replaced with a Rising Edge product, the Intelligent Display Controller (IDC). The IDC provides many of the

features of the original design with support for better LCD's, expanded I/O, and better tolerance to system failures.

With all the design changes complete, the second generation ORAID ADM was ready for delivery. In terms of system performance goals, the ORAID ADM exceeds the original program goal of 6 GBytes by 99% by offering 11.5 GBytes of storage. Performance factors were also exceeded by a wide margin. Burst transfer rates routinely occurred in the 12 MB/s region for reads and 6 MB/s for writes. Sustained transfer rates for reads approach 7 MB/s while writes occur in the 3.5 MB/s range.

Section 2 Introduction

The development of the ORAID ADM is intended to provide the United States Air Force with a new mass data storage and retrieval system. The system design will be based on existing technology to reduce development costs and commercial devices to reduce life-cycle costs. The system is designed to combine the benefits afforded by state-of-the-art RAID systems with the benefits of optical disk storage.

2.1 Background

There is a growing trend within the Air Force intelligence community to collect, process, store and disseminate a large amount of digital data. Advanced sensor systems are adding to the problem by providing data at higher resolutions on a 24-hour a day basis. In addition, the user community is demanding faster access to multiple data types, (e.g., imagery, maps, video, voice, and text). Intelligence products in the future will increasingly combine these data types in new and innovative ways to provide information to the right decision-maker at the right time. There is also a move from a centralized to a more distributed data processing environment, which will require mass storage system to interface and inter-operate with many different user workstations and data processing systems. With the changing intelligence data processing requirements, there is a growing need for more and better mass data storage solutions. Many Air Force programs, such as Air Force Intelligence Network, (AFINTNET), Electronic Footlocker and unit-level mission planners, will benefit from new mass storage concepts. New mass storage systems like the ORAID will help to meet these processing goals.

2.2 Significance of ORAID ADM Development

In developing new concepts and systems for mass storage, performance and reliability are undeniably important. The need for these features within the computing community are being met through the use of Redundant Array of Independent Disk (RAID) storage systems. Storage devices and technologies also continue to rapidly advance. When reviewing the potential storage devices that might be applied to new mass storage system development, the continued improvements in optical storage technology make this media a viable candidate for consideration. When a RAID architecture is considered in concert with optical disk technology, several reasons for merging these two dynamic technologies into a single unit begin to emerge. Table 2-1 details the benefits of each technology that when combined, make ORAID an important new concept in mass storage systems.

Table 2-1. Benefits of Optical Storage and RAID Systems

Benefits of Optical Disk Storage	Benefits of RAID Storage
High Capacity per Disk (Currently 2.6 GB)	Data I/O More Closely Matched to Host I/O
High Reliability ($> 10^7$ Rewrites)	High Reliability Through Redundancy
Transportable Storage	Parity Data to Prevent Data Corruption
Low Cost per MByte	Simplify Storage Management
No Media/Head Interface Problems	
Infinitely Expandable Capacity	

2.3 Development Goals

The ORAID ADM development focused on satisfying several basic goals. Among them

- Serve as an integrated storage product consisting of multiple optical disk drives designed to appear as a single disk to a host requesting/transmitting data. Optical drives will be used due to their high reliability, data retention capabilities and the ability to remove media for storage. The number and type of data drives shall be selected to support the performance goals described in Section 2.3.4.
- Support a RAID architecture to increase reliability and accuracy of the stored data. A standard RAID level(s) will be supported. The RAID level will be chosen to provide the best performance relative to the anticipated use of the system. Parity information will be retained in the RAID to reconstruct data in the event a single drive fails. In the event more than one drive fails simultaneously, the data cannot be reconstructed.
- Design the system to be modular and scaleable to take advantage of improved capabilities that may be available in the future.

The ORAID was not designed to serve as a real-time data collection device. All data to be stored on the ORAID is transferred from a host processing platform. Data retrieval is performed through the host on a file name basis. The host is responsible for all file management. The ORAID operates either as a dedicated resource to the host or as a multi-host resource through a file server. All connections between the controlling host and the ORAID are made through an industry standard SCSI interface.

2.3.1 Design Approach

In order to meet the development goals of the ORAID ADM, Rising Edge Technologies evaluated four possible approaches to meeting the technical requirements of the SOW. Each alternative is based on utilizing commercially available components to perform the final system implementation. One of the basic assumptions made for this analysis was that

multiple RAID control devices and or optical disk drives would be needed to achieve the rate and capacity requirements of the contract.

- Approach 1: Software Multiplexer. This method creates a software multiplexer that would drive two interface/control cards in the host workstation allowing access two independent RAID controllers. The software would be responsible for all controller and data management. Controller management would involve set up control of each RAID device as well as error handling and device synchronization (start, stop, logical sectors, etc.). Data management would be responsible for splitting the data to each controller, maintaining its storage locations, and reconstructing data into its original data sequence.
- Approach 2: Hardware Multiplexer. This approach would also rely on two RAID devices but utilize a hardware Multiplexer to perform data division to the two units. The Multiplexer would be responsible for buffering and flow control to each RAID device. This unit would provide a single interface to the host workstation.
- Approach 3: Extended Logical Unit. Extended Logical Unit (ELU) would utilize the 'tiers' available on most RAID controllers and rely on system software to select the next tier for data storage. A tier consists of a number of disk drives connected in parallel to the RAID controller. Thus, data is alternated between tiers to increase the storage capabilities of a single system. By implementing tiers, the system could be designed to meet the storage requirement, however, little would be gained in terms of data transfer performance.
- Approach 4: Integration of COTS. Most commercially available RAID controllers do not provide the parallel disk support to meet the depth of storage requirements of the SOW. The COTS approach would relax the SOW requirements to match what could be assembled based on the rate and capacity offered by a single, standard, commercially available RAID controller solution.

With the basic design approaches identified, the system was further decomposed into major subsystems. Each subsystem was then defined in order to conduct a thorough analysis and evaluation of candidate design methods and available components.

2.3.1.1 Optical Disk Storage Subsystem

The optical drives will support removable media. The drives shall be either full- or half-height and not exceed 5 1/4-inches in width. The drives shall be capable of supporting a standard, single-ended or differential Fast SCSI interface. Each drive shall contain a minimum of 512 kB of cache memory.

The first generation ORAID ADM was designed to utilize the current state of the art optical storage at the time of contract award. This consisted of storage media comprised of a removable cartridge with the capacity to store a minimum of 650 MBytes per side for a total storage capacity of 1.3 Gigabytes. The disk media was chosen to conform with one of the industry standards for optical media. Media selected is available from at least two independent sources.

2.3.1.2 RAID Controller

The RAID controller was purchased from a commercial source. Although this does not represent the most efficient or effective way to implement the ORAID ADM, it was viewed as the most expedient way to verify program goals. Modifications to a commercial vendors product were acceptable in order to meet the requirements of the ORAID ADM SOW. The RAID controller performs several basic functions. The selected controller met the following characteristics:

- serve as the interface between the host control platform and the disk array
- support the selected RAID Level
- provide a disk level interface to the selected optical drives
- provide a visual means for monitoring status of the ORAID

The RAID controller must maintain as high a data rate as possible between the ORAID and the host. This is accomplished through the use of a SCSI-2 Fast and Wide interface. This interface will support data transfer rates up to 10 MBytes/second. The interface between the optical drives and the RAID controller must be compatible with single-ended or differential SCSI Fast standards. Specific interface requirements will be determined by the final selection of the components.

2.3.1.3 Equipment Chassis

The equipment chassis is used to house the optical disk drives, RAID controller, cooling fans and power supplies. The chassis must have sufficient space to mount the RAID controller and all required disks with full front panel access. The ORAID chassis provides the capability to remove or insert the optical drive unit without opening or otherwise altering the physical state of the chassis.

2.3.2 Host Control Platform

2.3.2.1 Host Processing Platform and Operating System

The host processing platform and operating system have the greatest impact on the ability of the user to utilize the large data storage capabilities of the ORAID. The host workstation delivered as the primary array controller with the ORAID ADM is a Sun SPARCstation 5 running Sun Solaris 2.3. System support for other processing platforms

is also supported. Any recommendations for the workstation and operating system that would improve ORAID performance were also considered.

2.3.2.2 Host Data Interface

Data transfer to and from the ORAID is accomplished through a standard SCSI-2 Fast and Wide interface. All drivers required to operate this interface in the host processing platform were provided with the system.

2.3.3 Modes of Operation

The ORAID system is designed to support two basic modes of operation. A brief description of each mode is provided in the following sections.

2.3.3.1 Dedicated to a workstation.

All testing of the ORAID ADM was performed in this dedicated mode using a Sun SPARCstation 5 running Sun Solaris version 2.3. The ORAID should be backward compatible with other Sun operating systems. Dedicated operation will be performed by having the end user working directly on the Workstation connected to the ORAID.

2.3.3.2 File server on a client/server network.

File server operation will allow users connected to the host processing platform via a network interface to retrieve data files. Although no development is required to support this mode of operation, the ORAID was designed so as to not preclude this type of access.

2.3.3.3 System Software

The host processing platform will interface to the ORAID thorough a Fast and Wide SCSI-2 compatible interface. The operating system must support a standard SCSI-2 driver. The operating system selected for use with the host platform will perform all file management functions. In the event the operating system cannot support the storage capabilities of the ORAID, a file structure will be created that will support a minimum file size of 150 MBytes and a maximum size as large as possible for the selected operating system.

2.3.4 Performance Goals

The system shall be designed to meet the following performance goals:

- Storage Capacity: 6 Gigabytes
- Data Transfer Speed: 50 Mbits/sec (burst), 25 Mbits/sec (sustained)

Although additional features of RAID are desirable (i.e. component redundancy, hot swappable drives, etc.), these two performance goals are the most important features of

the ORAID. Trade-offs against other RAID features are acceptable if they help reach the desired performance goals. For example, traditional RAID systems do not have the ability to change the physical storage media (i.e. optical platter), therefore the ability to remove the drive while the RAID unit is powered on is important. Replaceable media may reduce the importance of this feature.

2.3.5 Interface Specifications

The ORAID supports an industry standard interface between the host processing platform and the ORAID and within all internal bus structures. No custom or proprietary interfaces have been used. At a minimum, the ORAID was designed to support a high speed data/control bus to the host processing platform and a data/control interface to each of the optical disk drives. The drive interface was viewed as a single bus or individual cables wired to each drive. The preferred interface standards are SCSI-2 Fast and Wide to the host processing platform and single-ended or differential SCSI Fast to the disks.

2.3.6 Maintenance

The ORAID was designed to provide maintenance features that allow for easy removal and replacement of critical system components and meaningful diagnostic routines that display overall status of the system. A user interface and audible alarms were provided for control and monitoring of diagnostics. Diagnostics include such basic features as media checks, disk status, and audible alarms for critical failures.

2.4 Contract Scope

Contract F30602-94-R-00446, Optical Redundant Array of Inexpensive Drives (ORAID) is intended to create an Advanced Development Model (ADM) of an optical disk-based RAID storage system. With the current advances in optical data storage and redundant systems, the basic development philosophy supports integration of existing components versus new design. The system is designed to make maximum use of Commercial Off The Shelf (COTS) components. System development evolved through several basic stages. Analysis was conducted to select the supported RAID Level, candidate components for the optical disk drive and supporting media, and the system RAID controller. The system utilizes several of the standard RAID Level classifications. Once the analysis was completed, selected components chosen to undergo evaluation. The evaluation phase included a review of functionality versus system goals, performance verification, software/hardware compatibility, and modification analysis. Modification analysis was performed to determine if changes were required to the COTS components to help them meet system performance goals. Upon completion of the evaluations, system components were ordered and modification/integration began. The final stage involved testing, documentation and delivery of a complete system.

The development of the ORAID followed a straight-line path from research and analysis through design to final testing. Each step in this process was described in general terms as

a phase. Table 2-2 lists the 5-phase process that was implemented and provides a general description of the goals intended for each phase.

Table 2-2. ORAID Program Development Model

Phase	Description
Analysis/Unit Testing	Research problem, gather information, develop preliminary architecture
Development	Develop detailed design and strategy
Design	New component development; Modifications to existing hardware/software;
System Integration	Integrate system components
Testing	Test system performance

This report will follow this development methodology and describe the results obtained during each phase.

Section 3 Analysis and Evaluation

The design of the ORAID ADM proceeded in two stages. The following section details the design, analysis and evaluation process used to deliver the first generation unit. Similar techniques were applied to the second generation unit as well. Additional details in the second generation unit will be provided in subsequent sections.

3.1 Analysis of Commercial Optical Disk Drives

One of the primary goals of the ORAID development was to utilize commercially available optical drives for data storage. In order to select the best drives to support system requirements, a thorough review of commercial drive vendors was performed. A list of manufacturers was developed through research and referrals. This list was then used to gather technical information so that product offerings could be evaluated. The goal of this review was to obtain a representation of the strongest candidate optical drives available on the market. Evaluation units would then be obtained for detailed performance analysis. Product offerings would then be compared to a set of weighted evaluation criteria in order to select the best drive for the system. The selected drive would later be combined with a similar list of RAID controller candidates. The various combinations of these candidate components would then be evaluated to determine the best overall system architecture.

3.1.1 Review of Optical Drive Manufacturers

In order to determine potential manufacturers for the optical drives, the type of drive had to be selected. A minimum set of performance criteria was established. The drive had to support a media that was rewritable, removable, and had a capacity of 650 MB (minimum) per side. In reviewing drive technologies,^{1,2} it became apparent that Magneto-Optical (MO) drives provided the required capabilities.

A list of drive manufacturers was developed through literature review and referrals. This list represents a sample of the total optical drive market. In addition to specific manufacturers, several optical drive manufacturer representatives were contacted as possible alternative sources of information. Each contact was asked for products that satisfied several basic requirements. Among these were:

Storage capacity:	A minimum of 1.3 GBytes of total storage
Transfer Rate:	Read/write rates approaching 2 MBytes/Second
SCSI Compatibility:	Either SCSI-1 or SCSI-2 compliant command set
Removable media:	Media available from third party sources; support use in jukebox applications

¹¹ Asthana, Praveen, "A Long Road to Overnight Success", IEEE Spectrum, pp 62-63, October 1994.

² Stowe, David, "The Optical Challenge", AdvancedSystems, pp 86-90, May 1994.

The resulting recommendations were tabulated. Table 3-1 provides a list of vendors contacted, product offerings, and a few notes associated with their product offering. Additional technical details were received from each vendor to aid in product analysis. Discussions with manufacturers' representatives did not prove to be as successful a method for obtaining information as initially anticipated. Therefore, all efforts in determining candidate products and technical specifications focused on working directly with the manufacturers.

Table 3-1. Optical Disk Drive Manufacturers

Manufacturer	Product	Notes
Fujitsu, San Jose, CA		Only optical is 230MB, 3.5-inch
Hewlett-Packard Greeley, CO	HP C1716T	Very slow transfer rates
Hitachi	OD152	1024 or 512 sector supported, 2GB drive
IBM, San Jose, CA	0632 Model CHA	1024 or 512 media, up to 2.28 MB/s
Laser Magnetic Storage Inc. (LMSI) by Phillips		12 or 5 1/4-inch; can't get rate
Maxoptics, San Jose, CA	T3-1300	max sustained read 2.2 MBs, write 1.1 MBs. Check O/S for sector size compatibility
MOST, Cypress, CA		Nothing in the storage range. Check IBM, Sony
Olympus Image Systems, Irvine, CA	Deltis series	Packaged as system for MAC or PC; all parts included (disk, media, controller, etc.)
Pinnacle Micro Irvine, CA	Sierra 1.3GB	High RPM (4500) yields fast transfers (2MBs), 1.3G drives
Sony San Jose, CA	3241 (internal) 3251 (external)	Burst 3 MBs, 1-2 MBs sustained. Media dependent, 1024 sector yields 1.3G disk

With the technical information obtained from the manufacturers, a careful analysis of storage capacity, transfer rates, SCSI compatibility, and form factor was conducted. It was hoped that this analysis would produce 3 to 5 sources of potential optical disk drive products.

3.1.1.1 Advantages of SCSI Compatibility

SCSI compatibility was selected as an important criteria for drive selection for several reasons. The SCSI standard, at the time development began, represented the fastest data/disk interface available. Thus the best transfer rate performance was expected to be available with SCSI compliant drives. There are many vendors providing SCSI compliant

RAID systems, therefore practically any commercially available RAID controller is a candidate for the system. In addition, the rapid development of higher density, higher performance SCSI compatible magneto-optical drives³ ensures a strong growth path for next generation ORAID systems.

3.1.2 Analysis of Optical Disk Drive Technical Specifications

The selection criteria for the disk drives was based on requirements outlined in the contract Statement of Work (SOW) and the results of user requirements analysis. A summary of essential requirements is provided in Table 3-2.

Table 3-2. Summary of Optical Drive Requirements

Criterion	Description	Available Range
Capacity	Capacity will be used to define the maximum amount of storage available on the optical media	230 MB to 2 GB
Transfer Rate	This will define the maximum sustained data transfer rate supported by the disk (i.e. disk transfer rate). It will not include information pertaining to the SCSI bus or any caching algorithm used.	947 kB/second to 3.02 MB/second
SCSI Compatibility	Specifically, compatibility with the SCSI-2 bus. SCSI-2 will be used to obtain the highest possible throughput to the host.	20 MB/second (Fast and Wide, burst)
Media	Will define removable media, availability from multiple vendors, standards, and format.	Format includes CAV and ZCAV, standards as applicable
Bytes/ Sector	The amount of data stored in a sector on the optical media.	512 or 1024 bytes per sector

With the large number of drives and variations in specifications, it became clear that selecting a single drive that met all desired analysis goals would be difficult. Therefore, two requirements were selected as the basis for comparison. The critical selection criteria used to determine candidate optical drives were maximum transfer rate and capacity. These two criteria were chosen over others since it is the user's goal to maximize system performance above other features. These criteria were then traded off against one another. Several of the listed parameters have a direct effect on each other. For example, the number of bytes per sector effects the amount of data that can be stored on the optical media. The number of bytes per sector (among other factors) will effect the disk data transfer rate.

In order to calculate the number of optical drives the system would require, several specifications were utilized from the SOW. The goal for total storage capacity of the

³ Costlow, Terry, "M-O unit aims at desktop", Electronic Engineering Times, pg 18, March 13, 1995.

ORAIID was established as 6 GB. The goal for transfer rate was established as 6.25 MBytes/second (burst) and 3.125 MBytes/second (sustained).

From the vendors contacted, drives offering the highest storage capacity and transfer rate were selected for comparison. Table 3-3 summarizes information for each drive selected. The capacity shown is indicative of the entire disk. Since only one side of the media is available for storage at any given time, half of the total storage available was used to determine number of disks required. The last column of the table, 'Number', indicates the number of drives that would be required to meet the transfer rate (TR) and storage capacity (SC) goals of the ORAIID system. In all calculations of transfer rate, the lowest sustained transfer rate of the drive specified by the drive manufacturer was used. These numbers were then used to help select COTS RAID controllers or devices for the final system architecture design.

Table 3-3. Summary of Selected Optical Drives

Manufacturer	Model	Capacity	Transfer Rate	Number
Hitachi	OD152	1.739 GB (512)	1300-2690 kB/sec	5 (TR, 512)
		2.023 GB (1024)	1540-3020 kB/sec	8 (SC, 512)
				5 (TR, 1024)
				6 (SC, 1024)
IBM	0632 Model CHA	1.19 GB (512)	912-2280 kB/sec	7 (TR, 512)
		1.3 GB (1024)		12 (SC, 512)
				6 (TR, 1024)
				10 (SC, 1024)
Maxoptics	T3-1300	1.19 GB (512)	1100-2200 kB/sec	6 (TR, 512/1024)
		1.3 GB (1024)		12 (SC, 512)
				10 (SC, 1024)
Pinnacle	Sierra	1.191 GB (512)	2000 kB/sec	4 (TR)
		1.305 GB (1024)		12 (SC, 512)
				10 (SC, 1024)
Sony	SMO-S521	1.2 GB (512)	940-1800 kB/sec	7 (TR, 512)
		1.3 GB (1024)	1000-2000 kB/sec	12 (SC, 512)
				7 (TR, 1024)
				10 (SC, 1024)

3.1.3 Discussion of Specific Optical Drive Parameters

3.1.3.1 Properties Effecting Data Transfer Rate

There are two formatting methods currently used for storing data on MO drives. The format method used has a direct effect on the transfer rates for a given optical drive. These formats are Constant Angular Velocity (CAV) and Zoned Constant Angular Velocity (ZCAV). CAV disks maintain a constant number of sectors per track. While this maintains a constant transfer rate, it utilizes far less than the potential capacity of the disk. ZCAV formatting allows the number of sectors per track to vary. The farther the track is

from the center of the disk, the more sectors it is capable of storing. Therefore, the amount of data stored increases as you move away from the center of the disk. Since most drives spin at a constant rate, data linear speeds vary from zone to zone. Thus data transfer rates tend to be slower for data stored in the center of the disk and faster for data stored at the edges.⁴

Another factor effecting disk performance is caching. Caching involves the use of Random Access Memory (RAM) for temporary storage of data. Reads and writes from/to cache can be performed at the SCSI transfer rate (less command overhead). When a disk reads or write is required, seek time, latency, disk transfer rate (determined by where the data is stored on the media), and SCSI command overhead all become a factor in determining the throughput. All of these factors have a significant impact on SCSI transfer rate. Although it is clear that cache can provide significant advantages for small transfers (i.e. files less than or equal to the cache size), ORAID will typically operate with large data files (100 to 150 MB). It may not be practical, given memory costs, to install such large amounts of cache in the ORAID controller. Therefore a careful performance trade study was performed to determine if cache should be supported and the optimal amount to maintain system performance. In addition, the RAID level implemented will also have a significant impact on performance (see Section 3.2).

The final property considered during drive evaluation was the read/write functions of the drives. Data reads are typically performed in a single pass. The drive will seek to the proper sector and return the requested data. Data writes however, are not quite as efficient. MO drives perform writes in a three step process of erase/write/verify. This process is very time consuming. Coupled with the inherent seek times, write performance was expected to at least 1/3-less than read performance. There may be options available from the drive manufacturers to reduce the time it takes to perform a data write, thus further increasing the data transfer rate.

3.1.3.2 Media Standards

There are several organizations currently certifying optical storage media. They are the European Computer Manufacturing Association (ECMA), the International Standards Organization, (ISO) and the American National Standards Institute (ANSI). All three standards bodies have approved the CAV and ZCAV formats. Although the basic concept is the same, the storage media is not compatible across drives supporting the different standards. Therefore, when selecting an optical drive, the supported media should conform to one of the recognized standards and be available from multiple sources.

⁴ Williams, Don, "A New 2-Gigabyte Data Storage Standard for Magneto-Optical Drives and Libraries", Hitachi America, Ltd., 1994.

3.1.3.3 System Requirements and Drive Performance

The primary factor considered in development of the ORAID was system performance. This not only includes transfer rate and storage parameters, but the ability for the system to operate effectively in the user's environment. Drives with high MTBF and low MTTR characteristics are desirable.

3.1.4 Testing Goals

A theoretical analysis was performed on each drive prior to selection. The primary intent of this analysis was to determine the raw performance characteristics of each drive being considered. Physical testing was used to validate the theoretical numbers and determine several other factors. Table 3-4 identifies the performance goals confirmed during drive testing.

Table 3-4. Optical Drive Performance Tests

Performance Test	Description	Sub-Categories
Data Write	Data Write performance will be measured for each zone, maximum, minimum, average, and streaming rates.	3 Pass Write 2 Pass Write Cache EN/DIS Media type
Data Read	Data Read performance will be measured for each zone, maximum, minimum, average, and streaming rates.	Cache EN/DIS Media type
Error Handling	Error handling performance will be tested to determine the drive response to multiple types of errors.	ECC EN/DIS Command Errors Media Errors
SCSI Compliance	Verify compatible with SCSI command set. This will help determine the effective changes required for OEM RAID controllers.	
Operating System Interface	Operating System testing will test for maximum supported file size as well as the effects of 512 and 1024 byte media on OS parameters.	

The testing was conducted on a PC-platform as well as a Sun SPARCstation. This dual testing path was selected to allow Rising Edge to use existing PC-based code and drivers to quickly characterize performance before developing any test routines for the SPARCstation.

3.1.5 Test Methodology

Performance testing was designed to measure the data transfer rate for read and write functions, compatibility of the disk interface to SCSI specifications, and physical performance characteristics of the drive such as spin-up time, re-calibration, and shutdown time. The test sequence involved determining the average transfer rate for reads and

writes, testing read and write performance for different zones on the disk, and observing the effects of cache on read and write performance. A description of each test is provided in the following sections.

3.1.5.1 Average Disk Reads and Writes

Full disk read and write testing consists of sequentially reading or writing data to or from the disk beginning with sector zero and ending with the last accessible sector on the disk. After the data transfer occurs, the period of time required to complete the operation is measured. Average transfer rate is then computed by dividing the total number of bytes read/written by the measured time value. This test is repeated for various block sizes in order to see the effects caused by SCSI overhead or the amount of time consumed performing I/O requests. The resulting data is then plotted for each block size.

3.1.5.2 Zone Reads and Writes

Zone testing is similar to full disk read/write testing. Data is read from or written to every sector on the disk; however, in zone testing, transfer time is measured for each of the sixteen zones on the ZCAV media. The resulting transfer rate of each zone is then plotted.

3.1.5.3 Cache Testing

Cache testing is performed by using burst data transfers to the drive. Transfer sizes ranging from $\frac{1}{2}$ the size of the cache to $1\frac{1}{2}$ times the size of the cache are used. The purpose of this test is to determine the effects disk cache can have on drive performance. Cache testing also helps clarify the maximum burst transfer rate of the drive. This is useful in determining performance of the ORAID in transactional processing environments. Cache testing is also used in conjunction with read and write performance testing.

3.1.5.4 Parity Tests

On drives that support parity testing, tests were conducted to determine the effects of parity on write or read transfers. Two types of parity are possible, SCSI bus parity or drive-based parity. Only drive-based parity was tested at this time. It is believed that the inherent parity protection offered by a RAID implementation will be sufficient to correct all data errors.

3.1.5.5 Test Matrix

Table 3-5 represents a simplified drive test matrix. Each test is performed with cache enabled and disabled. Once all tests are complete, data is entered into a spreadsheet and results are generated for comparison.

Table 3-5. Drive Test Matrix

Test	Cache Enabled	Cache Disabled
Average Transfer (Read)		
Average Transfer (Write)		
Zone Transfer (Read)		
Zone Transfer (Write)		
Cache Transfer		
Data Parity		

3.2 Analysis of Standard RAID Levels

Rising Edge Technologies completed a review of available RAID architectures in order to determine the best solution for meeting the contract requirements. Standard RAID is currently defined in 8 levels. The implementation of each level is designed for optimum performance given a specific application. This cursory review was performed to select and validate a RAID method. With a method selected, additional research was then directed to collect data on system components that would support the selected RAID level.

The concept of RAID and its various levels was presented in 1987 by Patterson⁵. The initial proposal defined RAID levels 0 through 5. Several RAID vendors have added their own proprietary levels as well as the additional levels added by the RAID Advisory Board⁶ (St. Peter, MN). RAID 7^{TM7-8} is a patented design architecture of Storage Computer (Nashua, NH) while RAID Levels 6, 0&1, and 53 are defined by the RAID Advisory Board. A description of these additional levels are included for completeness. These RAID levels are summarized in Table 3-6.

Table 3-6. Summary of RAID Levels

RAID Level	Description
RAID Level 0	Block Striping: RAID level 0 implements a data striping approach across all drives in the system. Blocks of data are stored sequentially on each drive in the system. RAID level 0 has no redundancy. If one drive fails, all data stored on the system becomes unusable. Due to its lack of fault tolerance, RAID level 0 is best suited for high performance applications that do not require data protection.

⁵ Patterson, D. , G. Gibson, and R. Katz, "A case for Redundant Arrays of Inexpensive Disks (RAID)", University of California, Berkeley, Report no. UCB/CSD/87/391, December 1987.

⁶ The RAID Advisory Board, The RAIDbook, A Source for Disk Array Technology, The RAID Advisory Board, Fourth Edition, September 1, 1994.

⁷ RAID 7 is a trademark of Storage Computer Corporation, Nashua, NH.

⁸ Obrien, John, "RAID 7 Architecture Features Asynchronous Data Transfers", Computer Technology Review, Winter 1991.

RAID Level 1	Disk Mirroring: RAID level 1 implements a mirror disk for each drive in the system. Use of a mirror drive virtually eliminates system interruption due to drive failure. Data is duplicated (in pairs) for storage on each drive. The use of dual drives more than doubles the cost of the storage unit. RAID level 1 provides fair read performance but much slower write performance since it must write to two disks.
RAID Level 2	Bit Interleaving: RAID level 2 interleaves data and error correction codes at the bit level across all drives in the array. The error correction code commonly used in the Hamming code. RAID 2 is typically used for large computer systems due to the robustness of its design.
RAID Level 3	Striped Array with Parity: RAID level 3 is similar to RAID level 2 but better suited to microcomputers. In a typical RAID level 3 system, only one parity drive is used for all data drives. Interleaving of data occurs at a predefined logical size (bit, byte, block, or other unit). Each drive holds a portion of the data. Synchronized spindles help yield high data transfer rates. This makes RAID level 3 ideal for high performance systems, such as image processing, where speed and accuracy are important.
RAID Level 4	Block Stripping with Dedicated Parity Drive: RAID level 4 is similar to RAID level 3. The primary difference is that RAID Level 4 systems operate drives independently while RAID Level 3 operates drives in parallel. Typically write operations are slow due to the parity read to a single drive that occurs before each write.
RAID Level 5	Block Striping with Distributed Parity: RAID level 5 contains both data and parity blocks on each disk, thus eliminating parity drives. Eliminating the parity drive allows the system to perform multiple read and write functions at the same time. This increases the virtual transfer rate by one-half the number of drives in the array, thus increasing performance as more drives are added.
RAID Level 6	Independent Parity: Two independent parity computations requires three drives to fail before data is lost. Similar in performance to Level 5 but much higher reliability.
RAID 7 TM	Embedded O/S and Asynch I/O: Complete I/O independence and asynchronous transfer to all connected drives. This is accomplished by three features: separate cache for each I/O interface; each interface is connected to a high speed data bus; a process oriented, real-time O/S is embedded in the array architecture.
RAID Level 0&1	A combination of RAID Levels 0 and 1. Provides RAID Level 0 performance with RAID Level 1 redundancy (disk mirroring). Provides significant gains in performance and reliability but at a high cost.
RAID Level 53	A combination of RAID Levels 0 and 3, Level 53 provides the data transfer performance of RAID 3 with the I/O request performance of RAID 5. A RAID 0 striping layer is used in front of the RAID 3 layer to yield a highly reliable storage system with RAID 3-like costs.

Rising Edge considered several factors as key to selecting a RAID level to support system design. These are:

- High data transfer rate for array reads and writes
- Error correction through the use of parity

- Ability to support large, single files
- Available from a commercial source

The first three factors will be satisfied by a direct analysis of the RAID level. The selection of a specific level then allows additional research to be performed to determine commercial support. An analysis of the first three topics against the seven basic RAID levels is presented in Table 3-7.

Table 3-7. Analysis of RAID Features

RAID Level	Transfer Rate	Error Correction	Large Files
RAID Level 0	I/O transfer rate increased as multiple of drives.	None. Single drive failure can result in lost data.	Supported.
RAID Level 1	I/O is one-half of total drives in system.	Mirroring provides two copies of all data. No real "error correction".	Supported
RAID Level 2	Driven by transfer rate of slowest disk.	Error correction interleaved across drives	Supported
RAID Level 3	Spindle sync yields high I/O	Single parity drive required for all drives installed.	Supported
RAID Level 4	Fast reads for small files.	Interleaved at the transfer level. Parity check on a file basis.	Supported but interleaved parity causes performance penalty for large files.
RAID Level 5	High overhead to track parity on all drives reduces throughput.	Parity stored across all data drives.	Best suited to small I/O requests.
RAID Level 6	Improved performance over Level 5	Independent parity calculations with separate I/O	Supported but may not be easy to achieve
RAID 7 TM	Limited to the maximum throughput of any single disk.	Independent parity drive with separate I/O	Design optimized to support many small sequential access requirements.

This first phase RAID analysis strongly supported the implementation of RAID Level 3. All of the general requirements of the ORAID contract are met: speed, error correction, and large file support. Although many of the other levels have features that meet some of the ORAID requirements, only Level 3 has all the features to support a full implementation of the ORAID. A Level 3 implementation can also grow in capability as drive speeds and storage capacities increase in the next few years. The next step is to determine the commercial support available for RAID Level 3 controllers and

components. In searching for commercial support, Rising Edge considered both RAID system manufactures and OEMs of components used in RAID systems. The next step was to determine if sufficient vendor support existed in commercially available RAID 3 products.

3.3 Analysis of Commercial RAID Controllers

Over 70 vendors of RAID systems and associated products were discovered⁹. However only a handful of companies manufactured the RAID controllers that made these systems possible. Since it was the intent of Rising Edge to utilize as much commercially available equipment and components as possible, an investigation of the controller manufacturers was conducted to determine the current level of products supporting ORAID-type applications. Rising Edge also investigated system manufacturers who might be interested in providing their proprietary controllers in an OEM version.

3.3.1 RAID Controllers

The concept for the development of the ORAID assumed a commercial RAID controller would be available to support system needs. There are many vendors supplying RAID systems to the commercial marketplace. Rising Edge was convinced that a vendor could be identified to provide components of their products to support the ORAID development. The RAID controller selected must meet some basic high level requirements. Namely:

- Support Level 3 of the standard RAID model definition
- Be supportable as a stand-alone product.

A significant number of RAID manufacturing companies were discovered. By comparison, relatively few OEM controllers manufacturers exist. Telephone conversations and a review of product specifications and literature was conducted. The results were then combined with the optical disk research to create a recommended system architecture.

3.3.2 Review of RAID Controller Vendors

A list of RAID system manufacturers and OEM controller manufacturers was created through literature review and referrals. The list was developed to represent a sample of the most influential players in the RAID market. In addition to specific manufacturers, several RAID manufacturers' representatives were contacted as possible alternative sources of information. Each contact was asked for products that satisfied several basic requirements. Among these were:

⁹ Spencer G., "RAIDing with high availability", Advanced Systems, September 1994, pps 55 - 68.

- Block/sector Support: Maximum number bytes per sector supported.
- Capacity: Maximum number of drives supported; future expansion.
- SCSI Compatibility: Either SCSI-1 or SCSI-2 compliant to the host as well as the drive.
- SCSI Command Set: Since all vendors traditionally support magnetic media, it was important to identify supported SCSI commands that may differ for optical disk drives.
- Firmware Support: Ability to alter performance of the controller via modifications to firmware.
- Hot swap: Inherent ability of the unit to support hot swap features.

In conversations with the commercial vendors, the common themes of user application and operating system were frequently discussed. Many of the available systems are tuned to perform in specific applications and with specific operating systems. This is due in large part to their customer base consisting primarily of client/server and database management applications. In describing the application intended for ORAID, most were unable to offer any advice on how their basic controller would operate.

The list of RAID systems offered by different manufacturers is extensive. In discussions with these vendors, it was discovered that most pre-packaged RAID solutions used an OEM controller provided by a third party source. Most of these manufacturers did not wish to reveal the source of their controller. However, research led us to several OEM sources who were contacted for additional information. Table 3-8 provides a small sampling of the list of the manufacturers contacted for information regarding their RAID controllers. These few were selected for further discussions based on their ability to provide their controller board(s) as an OEM product. The table includes several major OEM manufacturers as well as two system manufacturers willing to provide their controller as OEM products.

Table 3-8. RAID Controller Vendors

Manufacturer	Product	Notes
BusLogic Santa Clara, CA	DA-2988	PCI and EISA boards
Ciprico Plymouth, MN	6700/6800	RAID 3; Controller is part of their system; 8+1
CMD Technology Irvine, CA	CRD5000	RAID 0, 3, 5; OEM unit
Digi-Data Jessup, MD	Z-RAID	RAID 0, 3, 5; OEM unit

Mylex Fremont, CA	DAC960	EISA controller, RAID 0, 1, 5 (6, 7)
Storage Concepts Irvine, CA	Concept 910	RAID 3; controller for their 910 (8+1); Jan 94 release; will OEM
Ultrastore Irvine, CA	UltraArray	RAID 0, 1, 4, 5

3.3.3 Analysis of RAID Controllers

Each of the controller listed in Table 3-8 were examined against the basic requirements. A partial listing of the results is provided in Table 3-9. In all instances, SCSI compatibility and the availability of the SCSI command set were part of the product offering. However, three of the vendors could not meet the basic RAID Level 3 requirement. The remaining four were contacted for specific technical information regarding implementation of their controllers.

Table 3-9. Controller Product Specifications

Manufacturer	Product	RAID Level	Firmware Mods	Width
BusLogic Santa Clara, CA	DA-2988	0, 1, 5	N	Up to 4
CIPRICO Plymouth, MN	6700/6800	3	Y	4 Data, 1 Parity or 8 Data, 1 Parity
CMD Technology Irvine, CA	CRD5000	0, 3, 5	Y	Up to 6 Data, 1 Parity
Digi-Data Jessup, MD	Z-RAID	0, 3, 5	Y	4 Data, 1 Parity
Mylex Fremont, CA	DAC960	0, 1, 5 (6, 7)	N	Up to 3
Storage Concepts Irvine, CA	Concept 910	3	Y	8 Data, 1 Parity
Ultrastore Irvine, CA	UltraArray	0, 1, 4, 5	N	

3.3.4 Sector Size Compatibility


Optical media is available with 512 or 1024 bytes per sector. Since the current line of RAID controllers are designed to primarily operate with magnetic devices, they only support 512 bytes per sector. All vendors contacted indicated that there would be major changes to firmware required in order to support a 1024 byte sector. These modifications would take a significant amount of time and testing to complete. By utilizing optical media with 512 bytes per sector, the firmware available in existing controllers could be utilized in direct support of the optical drives.

3.3.5 Spindle Synchronization

Spindle synchronization was considered an important factor in allowing the ORAID to perform at a high data rate. In discussions with the controller vendors, spindle synchronization features were stressed as an important function for the requirements of the ORAID contract. Many of the RAID 3 compliant vendors indicated that spindle synchronization was supported by their product but not required when using the controller in RAID 3 mode. Figure 3-1 illustrates some of the basic principles employed when considering spindle synchronization.

Spindle synchronization is designed to allow all drives in the array to start on the same sector and maintain a constant vertical alignment of sectors across all drives. Thus, whenever a write or read command is issued, all heads are positioned over the correct sector and as soon as the disks rotate to that sector, data transfer can begin. Figure 3-1(a) illustrates a case of non-synchronized spindles. Figure (b) shows the effects of vertical synchronization. In theory, this vertical sector alignment allows the system to begin transferring data to all drives simultaneously. However, consider a typical media with imperfections that create bad sectors, as in Figure (c). When a bad sector is one of the elements in the vertical alignment of the array, data can not be written to that sector when a transfer is initiated. The data would either have to be buffered or all data would have to be delayed until a vertical alignment of sectors was found that had no bad sectors. Even if such a column of sectors were found, there would be no guarantee that a sector somewhere within the resulting data block would also be marked bad.

One possible solution to this bad sector dilemma would be to mark sectors bad across the entire vertical section. However, this approach tends to fragment the disk into sections with small maximum usable files. The maximum file size would vary from system to system since there would be no way to guarantee the same bad sectors across random samples of disks. In addition, the maximum file size stored would be determined by the smallest continuous storage segment available on the drives.



The preferred implementation is to provide each drive with its own cache memory. The cache would be used to temporarily hold data when a bad sector is encountered. The data would then be written to the drive when the next sector is available. This use of cache is the basic principal the RAID Level 3 controllers rely on when spindle synchronization is not used.

The advantage of cache in a Level 3 implementation opens up the possibility of using almost any available optical drive. Initially it was believed that the selected optical drive would have to be modified to support spindle synchronization. However, the caching approach allows optical drives to be used without modification.

3.3.6 System Requirements and RAID Controller Specifications

The RAID controller was selected in conjunction with the optical drives. The speed of the drives, capacity, and number of drives support by the controller were all factors in determining which controller was selected. The width of the controller also has a direct effect on the block request size implemented by the host. The selected drive determines the optimum block size that will achieve maximum data transfer. For an n -wide system, the host block size will be n times the disk block size. The total file size to be transferred will then be a multiple of the host request size.

The host interface is supported through a SCSI-2 connection. Most of the controllers reviewed support the fast and wide implementation. Fast and wide is the preferred host interface since it will support transfer rates well above the systems goal of 50 Mbps.

3.4 General System Requirements

3.4.1 ORAID Power Requirements

A power budget needed to be established to determine the load requirements of any selected system power supply. An evaluation unit of the selected power supply was ordered prior to purchasing the final system unit. This supply was used to operate the prototype system and verify its performance under multiple conditions of operation and failure. Once the supply was proven to perform satisfactorily, a second unit was purchased and the two were configured for redundant operation. For preliminary testing, the power budget was used to select bench supplies to power the test units.

Several power supply vendors were contacted to discuss the ORAID system requirements and determine if commercial units were available. There are many commercial power supply vendors, so to reduce the number considered, only vendors that offered removable or hot-swap and redundant supplies were called. Table 3-10 lists vendors who had power supplies meeting these requirements.

Table 3-10. Power Supply Vendors

Vendor	Model
HC Power	PS Series
Astec America, Inc.	(first quarter 96)
Kepco	HSP Series
Industrial Resources	HP6261B
Omega	HP Series
Global Power	MG or HP series
Condor	Custom

3.4.2 Thermal Considerations

The application notes provided with each optical drive were studied in order to determine any special air flow or cooling requirements. Although the required air flow is typically small (on the order of 3 cubic feet per minute per drive), it is critical for proper drive operation. Packaging of the ORAID will be important in order to maintain proper operational temperature.

Section 4 Results

The implementation of the first generation ORAID ADM grew from the analysis described in the previous section. The following sections describe the results of specific testing and the resulting implementation to create a final deliverable.

4.1 Results of Optical Drive Analysis

During the analysis phase of the ORAID contract, five optical drives were selected for evaluation. Drives were chosen for various performance considerations and represent an excellent cross section of available optical drive technology. The selected drives are listed in Table 4-1.

Table 4-1. Selected Optical Drive

Manufacturer	Optical Drive Product
Hewlett-Packard	C1716T
Hitachi America	OD152S-1
Maxoptics	T3-1300
Olympus Imaging Systems	Deltis
Pinnacle Micro Systems	Sierra OHD-1300

The initial design phase performed a detailed technical evaluation of each drive. The drives were ordered as evaluation units for testing to avoid committing any funds to a drive that may not be suitable for final system design. Drives were ordered in OEM versions, i.e. no chassis, power supply, or SCSI controller card. Rising Edge provided these external elements as part of unit testing.

4.1.1.1 Anticipated Performance

At the time of first generation development, three optical drive types were available on the market. Each type represents an evolution of the original optical recording methodology. The technologies are labeled as 1x, 2x, and 3x, representing the level of evolution from the original standard. In order to determine the best technology to use, a theoretical performance estimate was performed to select the candidate drives. Performance was calculated based on each drive's physical characteristics. Critical factors were the physical format of the media and the rotational speed of the drive. An estimation algorithm was used to estimate the number of bytes per track across all zones of the media. This number was combined with the rotational speed of the drive to determine transfer rate. Figure 4-1 through 4-3 illustrate the resulting performance curves.




Figure 4- 1. Calculated Transfer Rate, 1x Drive




Figure 4- 2. Calculated Transfer Rate, 2x Drive




Figure 4- 3. Calculated Transfer Rate, 3x Drive

4.1.2 Test Configuration

The basic test configuration used to measure actual drive performance is shown in Figure 4-4. The PC workstation utilized an Adaptec AHA-2842A VL-to-Fast SCSI host adapter. Testing with the SPARCstation was conducted with the standard SCSI interface provided with the workstation. For testing, the SCSI controller was set to bus address 0 and the drive under test to bus address 2. The controller and drive were set to provide active termination. The drive was set to provide the SCSI TERMPWR.




Figure 4- 4. Basic Disk Test Configuration

A specialized SCSI bus analyzer was used to measure all performance criteria. The Ancot DSC-216/FD SCSI Bus Analyzer was used to perform transfer and command analysis. The DSC-216 Event Trace memory and Passive Tracer functions captured commands and recorded traces to allow the timing to be analyzed. The Analyzer acted as a passive bus monitoring device on the SCSI chain. Termination and TERMPWR were provided by the drive under test.

The drive under test was exercised through all switch or software selectable functions. Of specific interest were any switch functions that effect data transfer performance. Examples of such functions include write-to-cache or write-through-cache and write/verify/erase (3 pass) vs. write and verify (2 pass) operation.

Test software was developed to allow full control of all disks under test. Specific commands were utilized to monitor drive performance. All testing was driven from a test program that allowed specific commands and data patterns to be sent the drive. Additional features such as transfer rate analysis were also employed. These functions helped fully characterize disk performance. All test results were recorded and compared against one another to determine the drive that best met the system performance requirements.

4.1.3 Optical Drive Testing

The drive under test was housed in a chassis external to the Workstation driving the tests. Power was obtained from the Workstation power supply through extension cables. The SCSI Bus analyzer was connected to monitor the commands and responses between the host controller (initiator) and the disk under test (target). All commands were monitored to determine the activity of the bus and perform timing of each data transfer.

Optical drives were tested using two different media types. The selected media was compatible with the Zoned-Constant Angular Velocity (ZCAV) format. This format was selected for its increased storage density and higher performance with respect to transfer rate. Testing utilized media with 512 bytes per sector as well as media with 1024 bytes per sector.

A sequence of tests was devised to evaluate the basic functions of the optical drives that are key to ORAID performance. An inspection of drive commands was performed to compare the drive's command set to the ANSI standard command set. Once it was determined that the desired commands were supported, drive testing began. A set-up and verification test was performed. This test included verification of media installation and size, drive locking, and reading of bad sector data. Testing then began to determine drive transfer performance. Performance testing was conducted to include various combinations of file and block sizes, write/read locations, and media types.

4.1.3.1 SCSI Command Compatibility

SCSI command compatibility was performed through a comparative evaluation of the drive documentation and the ANSI X3.131-1994 Small Computer System Interface - 2 (SCSI-2) Standard. Section 16 of the specification lists commands that are to be supported by optical media. These commands were compared with the control commands supported by the optical drives for ANSI specification compliance.

4.1.3.2 Initial Set-up and Verification

Each drive manufacturer uses different methods of selecting default operating conditions for their drives. These settings were adjusted as required to ensure that all drives were tested under the same conditions. Testing was performed to evaluate the drives under all possible switch or programmable settings. Table 4-2 lists several key parameters that were checked prior to testing. It was important that these settings were known prior to testing so that test results can be properly interpreted. Switch settings and software controls are listed with the specific test tables for each drive. A matrix of possible configurations and settings was utilized to fully test all scenarios. The settings were then recorded in the test table.

Table 4-2. Switch Selectable or Software Selectable Commands

Setting	Description
OFF/ON	Force/Inhibit verify on Write command is default mode
OFF/ON	Read Cache is disabled/enabled is default mode
OFF/ON	Write Cache disabled/enabled is default mode
OFF/ON	SCSI Data Bus Parity check is disabled/enabled
OFF/ON	SCSI Synchronous data transfer disabled/enabled

4.1.3.4 Set-up and Verification Test

Preliminary testing was performed to determine the ability of the Workstation and driver software to communicate with the optical drive and verify drive functionality. The set-up sequence allowed some basic commands to be executed for each drive and confirm information about the optical media being used. A listing of the default switch settings and an explanation of their effect on drive operation were recorded prior to testing. Switches were then varied according to the test matrix to account for all possible conditions. To begin each test, the appropriate optical media was inserted in the drive under test and a specific command sequence was performed. The final step of the process pre-erased the media in preparation for data. Table 4-3 lists the set-up and verify commands in their operational sequence.

Table 4-3. Set-Up and Verification Test Commands

Command Sequence	ANSI SCSI OP Code	Description
PREVENT ALLOW MEDIA REMOVAL	1Eh	Prevent/allow media removal from the logical unit
REZERO UNIT	01h	Set the logical unit to a specific state
READ CAPACITY	25h	Request information regarding the capacity of the logical unit
MEDIA SCAN	38h	Scan media for contiguous set of written or blank blocks
READ DEFECT DATA	37h or B7h	Requests media defect data
ERASE	2Ch or ACh	Erase a specified number of blocks

Once the set-up and verification testing were complete, drive performance testing began. Performance tests were conducted on individual drives for all possible modes of operation.

4.1.4 Write/Read Performance

Read/Write performance testing was conducted to compare the actual drive performance numbers against calculated values. Performance testing was conducted using the following commands:

Table 4-4. SCSI Commands for Performance Testing

Command	ANSI SCSI OP Code	Description
READ (6)	08h	Transfer data from the media
READ (10)	28h	
READ (12)	A8h	
READ LONG	3Eh	Return data and ECC from media
WRITE (6)	0Ah	Write data to the media
WRITE (10)	2Ah	
WRITE (12)	AAh	
WRITE LONG	3Fh	Write data with ECC
WRITE AND VERIFY (10)	2Eh	Write data to media and verify
WRITE AND VERIFY (12)	AEh	
ERASE (10)	2Ch	Erase the specified number of blocks
ERASE (12)	ACh	

4.1.4.1 Data Write/Read Tests

Data write testing was used to determine the data transfer rate from initiator to target. Data read testing determined the data transfer rate from target to initiator. A matrix of possible test conditions was utilized to cover all possible test conditions for both the write and read data transfer scenarios.

Data written to an optical drive must undergo a two or three step process. These processes are typically write/verify or erase/write/verify. The selection of the process is indicated as a two or three pass write sequence in the test tables. Data write performance was also tested using the optical drive's internal cache. All optical drives offer a different amount of cache. By testing drives with different amounts of cache, it was possible to determine the effects cache size had on transfer rates. Tests were run using write caching as well as write-through caching.

Transfer rate was affected by several key factors. Table 4-5 lists each of these factors and an explanation of how they were expected to affect drive performance.

Table 4-5. Key Factors Affecting Drive Performance

Performance Factor	Description
File Size	File size will determine the number of sectors that are required for storage. The larger the number of sectors, the slower the average transfer rate for the entire file will become.
Block Size	The block size determines how much data is written to a single sector. This will affect transfer rate if the drive has to locate multiple or fragmented sectors to complete an entire transfer operation.
Data Location	Data located at the outer edges of the disk has a higher transfer rate than data located on the inner tracks of the disk.

For the purposes of testing, an assumption had to be made on the final drive configuration of the ORAID. The number of drives in the final configuration was used to determine the file size used in testing individual disk performance. In the early phases of system development, it was believed a COTS RAID controller would be discovered that would allow the ORAID to be configured as an 8+1 system. Assuming a typical file size of 150 MBytes, this would dictate each data drive hold a maximum of 18.75 MBytes per file.

In writing these files, several block sizes were tested. The smaller the block size the slower data transfer rate was expected to be. Optimal block size would have to be determined by the final selection of a RAID controller and the host operating system. In order to determine the impacts of data location, the disks were divided into several sections and file data would be stored to these sections. File transfer performance was then compared across all sections.

4.1.5 Tabulation and Recording

Each drive was tested individually. A detailed test matrix was used to distinguish differences between the default settings and configurations of each drive. Testing was conducted in accordance with this matrix and a standard test table was used to record results. Detailed test results for the five optical disk drives evaluated are provided in the Appendix. A brief analysis of results and a look at performance for drive reads and write are provided in the following sections.

With all raw test data in hand, results were then evaluated to determine the best performing drive. This drive was then further analyzed to determine “soft” factors such as vendor support, anticipated life cycle, availability, as well as other factors. Once all information had been collected, a final selection was made of the ORAID optical disk drive.

4.1.6 Results

4.1.6.1 Drive Testing: Olympus MOS525E

Test results were collected and compared with published specifications. Drive performance proved to be somewhat slower than the published specifications. Detailed test results of the drive performance are provided in the Appendix. The Olympus drive does not offer a read cache option. Therefore, only one read test was performed to determine average read transfer rate. The average read and write transfer rates are shown in Figures 4-5 and 4-6.




Figure 4- 5. Average Read Rate, Olympus MOS525E




Figure 4- 6. Average Write Rate, Olympus MOS525E

The Olympus is configurable for 2 or 3 pass write operations. Figure 4-6 illustrates the effects of each type of write transfer on overall transfer rate. Additional write cache tests were performed and are provided in the Appendix.

4.1.6.2 Drive Testing: Hewlett-Packard C1716T

The C1716T was added to the list of drives under test at the suggestion of a Maxoptics dealer. The drive was bypassed in early analysis due to its low RPM (2400 vs. 3000 or 3500) rate for spinning the media. However, with cache enabled, the drive provided respectable transfer rates for large block sizes. Figures 4-7 and 4-8 illustrate the average transfer rates for read and write data transfers.




Figure 4- 7. Average Read Rate, HPC1716T

The Hewlett-Packard C1716T offers a read cache option. The average read transfer rate was tested with cache enabled and disabled. The test results show cache has very little effect on the read performance of the drive. It is believed that the small cache buffer in the C1716T limits the performance gains that can be achieved through the use of read cache. A larger buffer would have a more dramatic affect on drive performance.




Figure 4- 8. Average Write Rate, HPC1716T

The C1716T operates in a 2 Pass write mode. The HP 2 Pass write is an erase/write process. No verify option is provided. Figure 4-8 shows significant increase in write performance cache provides for the HP write operation. Even though the C1716T has a small cache, a performance increase is achieved due to the slower speeds associated with data transfer during write operations.

4.1.6.3 Drive Testing: Hitachi OD-152-S1

The Hitachi OD152 is the only drive currently available supporting the 2 GB (3x) storage media. Discussions with competing optical disk manufacturers indicate the 2 GB media format has not been approved as a standard. A technical representative from Hitachi America indicated the ISO DIS for the 3x media is finished. The standard has been assigned number ISO/IECF 13842 and should be an accepted standard by Fall, 1995.

The lack of a standard for the 3x format caused some concern as to the future viability of the product. Hitachi indicates that the drive is selling well and future generations utilizing the 4x media are currently in development. All future drives will be able to read and write the current 3x media. Hitachi also indicates that the OD152 is being used extensively in jukebox applications by several large OEMs.

The OD152 had excellent data transfer performance for large block transfers. However, as with all of the other drives, this performance drops quickly as the block size is decreased. The reason for this decreases is the SCSI overhead associated with each transaction. The smaller the block, the more overhead associated with its transfer and the slower the data rate. Write performance suffers with the Hitachi due to its lack of a write cache. Figures 4-9 and 4-10 show the average read and write rates for the Hitachi optical drive.




Figure 4- 9. Average Read Rate, Hitachi OD152




Figure 4- 10. Average Write Rate, Hitachi OD152

4.1.6.4 Drive Testing: Maxoptics TMT3-1301

The Maxoptics T3-1300 was to be tested in its 1 MB cache and 4 MB cache configurations. However, the 4 MB unit was not available in an evaluation unit and could not be obtained until after the test cycle was scheduled for completion. Testing was therefore performed only on the 1 MB cache version. Although the T3 supports cache enabled and disabled functions, no discernible difference in drive performance was detected between the two modes.

It was discovered after the drive was returned that there was an error in the documentation. Although the drive supports software MODE selection of cache functions, this feature is overridden by a hardware jumper on the unit. During all tests, the jumper was set to enable caching. Thus, even though cache was disabled through software during testing, all cache tests were performed with cache enabled.

Figures 4-11 through 4-12 depict the average read and write transfer rates for the 1 MB cache Maxoptics T3 optical drive.




Figure 4- 11. Average Read Rate, Maxoptics T3




Figure 4- 12. Average Write Rate, Maxoptics T3

4.1.6.5 Drive Testing: Pinnacle Sierra

The Pinnacle Sierra was the last drive tested. The Sierra does not provide a means to disable the cache. Therefore all testing was performed with the 4 MB cache enabled. In order to find a way to gauge the performance of the Sierra drive, testing was performed with the bus disconnect enabled and disabled. The bus disconnect feature allows control of how much data is held in the disk buffer before the drive reconnects to the bus before a bus transfer can occur. Various buffer sizes were used with little effect in drive transfer rate.

As an operational note, the Pinnacle drive operated cooler than any other drive tested. Figures 4-13 and 4-14 show the average read and write rates for the Pinnacle Sierra.




Figure 4- 13. Average Read Rate, Pinnacle Sierra




Figure 4- 14. Average Write Rate, Pinnacle Sierra

4.1.7 Comparative Analysis of Tested Drives

The test results for each drive were then compared to determine the best drive for use in the ORAID. Table 4-6 illustrates the ranking criteria and comparison matrix used to select the ORAID optical drive.

Table 4-6. Drive Comparison Criteria.

	Olympus	HP	Maxoptics	Hitachi	Pinnacle
Ease of Use	High	High	Medium	Medium	Medium
Auto Spin Down	30 min; SCSI to start	30 min; SCSI to start	N/A	N/A	30 min; SCSI to start
Avg Write (kB/s)	167.6	374.7	363.5	229.7	389.2
Avg Read (kB/s)	526.4	389.6	624.3	780.5	517.5
Cost	\$1873	\$1900	\$1828	\$1950	\$2070
ISO Standards	ISO/IEC MA	ISO/ANSI/EC CMA	ISO/ANSI/E CMA	ISO/IECF 13842 (May 1995)	ISO/ANSI/E CMA

This table combines subjective and quantitative samples to yield an overall impression of the drive. Each factor was then weighted for importance to the ORAID ADM and a total score created. The weighting factor was based on a numerical scale from 1 to 10, with 10 as the most points available in any category.

In reviewing the preliminary table entries for each drive, no clear winner was established. It was assumed that one drive would have superior read/write performance over all others and thus be the logical choice. However, this proved not to be true. Although there was clear dominance by one drive over another, this dominance was related to the block size of the data being sent. Different drives performed better for large and small block sizes.

Therefore, it was necessary to perform additional testing in order to properly weight each result. Specifically, the drive performance in relation to the target operating system. In order to better understand how Sun/OS would effect system performance, testing was performed to determine how the Sun SCSI driver operated.

The SCSI bus analyzer was attached to the Sun. Typical bus activity was monitored during idle states and while data transfer and retrieval to SCSI devices occurred. This allowed some block size performance numbers to be collected. Typical transfers showed that the Sun SCSI driver used different block sizes for each transfer, and that the blocks tended to be large. The smallest block size detected was 8k bytes, while the largest approached 64k bytes.

The RAID controller will further divide these blocks based on the number of drives per system. Using the 8k block as an example, a system with 8 drives would write a 512 byte block to each disk. A 4 drive system would utilize a 1024 byte block to each drive. The large block (64k) would yield an 8k and 16k byte block, respectively.

Table 4-7 lists the resulting weighted scores for each drive. Drives with better small block performance received higher weights than those with outstanding large block performance. The drive receiving the highest score was then selected for use in the system.

Table 4-7. Weighted Drive Comparison.

Category (Weight)	Olympus	HP	Maxoptics	Hitachi	Pinnacle
Ease of Use (7)	9	10	7	6	7
Auto Spin Down (3)	6	6	10	10	6
Average Write (10)	4	7	9	6	10
Average Read (10)	7	5	8	10	8
Cost (6)	9	9	10	7	6
ISO Standards (4)	10	10	10	4	10
Storage Capacity (3)	7	7	7	10	7
TOTALS	306	323	370	320	344

The Maxoptics and Pinnacle Micro drives were ranked as the best drives of those tested. Their overall scores were very close, thus making the selection of one over the other difficult. In order to make a better determination, the transfer rate data from each drive was carefully reviewed. The Maxoptics drive appeared to have better rate balance between reads and writes. Based on the results of this analysis, the Maxoptics drives were selected for use in the ORAID system.

4.1.8 Optical Disk Drives: Lessons Learned

Very few problems were encountered while implementing and performing the tests on the candidate optical disk drives. Documentation provided with each test unit allowed for easy set-up of drive parameters to accomplish specific test functions. SCSI compatibility allowed for easy integration into a single control routine that was used to test all drives. The most noticeable discrepancy between expected performance and test results came in the area of specified transfer rates.

Several reasons were commonly discovered for the difference between specified and measured transfer rates.

- Use of the SCSI Specification: Many vendors publish the maximum synchronous and asynchronous data transfer rates of the SCSI interface within their specification. This number has very little to do with the actual performance of the disk and typically ignores the bandwidth lost to SCSI overhead (approximately 20% of the available bandwidth).
- Maximum Transfer Rate: Vendors also commonly publish the maximum transfer rate the drive is capable of providing, not sustaining. This can be seen on data sheets as specifications using phrases such as "up to .." in front of the transfer rate number. Due to the format of the optical media, data transfer rates tend to be faster near the outer edges of the media. Thus, the maximum transfer rate is only sustainable for a small portion of the disk. The rate decreases as the data is stored toward the center of the disk.

- Optimization for Data Transfer: It was discovered during testing that it is possible to select file sizes that optimize the transfer performance of the optical drives.

All testing designed by Rising Edge measured true sustained performance across the entire optical disk drive. It was believed that this type of test would provide the best representation of true drive performance.

With the rapid changes in optical drive technology, today's best choice could quickly turn into tomorrow's second best choice. By utilizing an industry standard device interface like SCSI-2, the growth path remains open for future versions of ORAID systems. Adopting the next generation of disk drives can be as simple as plugging the drive into the chassis and reconfiguring host parameters. Rising Edge will continue to monitor changes in optical disk drives technology in order to determine the impact they will have on the ORAID.

4.2 Results of RAID Controller Analysis

Completion of drive testing was important for moving forward with RAID controller testing. The controllers ordered for testing are typically shipped without disk drives. Therefore, in order to fully test the capabilities of the controllers, disk drives must be attached. Initially, a mix of hard disk drives and optical disk drives were to be used. However, the timing difference between the hard drives and optical drives would cause the controllers to fail the slower drives (the optical drives), and thus remain in a non-functional mode. Thus the minimal number of optical drives needed to keep the controller operational were installed for testing.

4.2.1 RAID Controller Vendors

The survey of RAID vendors performed during the analysis phase of the contract uncovered only two vendors (Ciprico and Storage Concepts) capable of providing an 8+1, RAID 3 compliant controller. However, Ciprico decided they were not interested in supporting Rising Edge in the development of an ORAID. Their opinion was not to stop production on their current products in order support the changes their controller would require to operate with ORAID. The withdrawal of Ciprico from the list of possible candidates left the list of potential commercial RAID controllers unrealistically small. In an effort to gain a reasonable market sample for evaluating RAID controllers, new vendors were added to the list.

The search for replacement controllers began by contacting magnetic disk-based RAID manufacturers. Of those contacted, several indicated they were purchasing their controllers from Symbios Logic. Symbios Logic, formally NCR, produces an OEM RAID controller line consisting of several products. Their ADP93 product offers RAID 3 support in a 4+1 configuration. A Symbios application engineer indicated a strong desire to support Rising Edge in the development of an ORAID controller. Their desire to

support the development quickly ended when Symbios announced their plans to discontinue the ADP93 in October of 1995. Symbios currently has no other plans to develop an OEM RAID 3 compliant controller.

Digi-Data Corporation was contacted during the analysis phase of the contract but ruled out due to the size (4+1) of their controller. However, with so few controllers supporting large numbers of disks available, the possibility of having to utilize a smaller controller became a reality. Thus, Digi-Data was reconsidered.

The CMD Technology CRD-5000 was the first candidate RAID controller identified for the ORAID project. The CRD-5000 is probably the most widely used OEM controller in the RAID market today. The unit can be configured as a 6+1 RAID 3 controller. Although CMD expressed no interest in providing technical support to modify the controller, testing was performed to see how an unmodified unit would perform with the optical drives.

As a result of the COTS RAID controller search, three RAID controllers were selected for testing. The Digi-Data, Storage Concepts and CMD controllers were tested. Table 4-8 presents a summary of the salient features of each controller. Preliminary testing was performed to establish compatibility of the base controller with the optical drives. The test includes the ability to set up the controller for different media, transfer rates, and spin up times, as well as the ability of the controller to recognize the optical drives on initial power up.

Table 4-8. Controller Product Specifications

Manufacturer	Product	RAID Level	Firmware Mods	Width
CMD Technology Irvine, CA	CRD-5000	0, 3, 5	Y	Up to 6 Data, 1 Parity
Digi-Data Jessup, MD	Z-RAID	0, 3, 5	Y	4 Data, 1 Parity
Storage Concepts Irvine, CA	Concept 910	3	Y	8 Data, 1 Parity
	Concept 810	3	Y	4 Data, 1 Parity

4.2.2 RAID Controller Testing

Preliminary testing was first conducted to establish compatibility of the base controller with the optical drives. This test includes the ability to set up the controller for different media, transfer rates, and spin up times, as well as the ability of the controller to recognize the optical drives on initial power up. All of the controllers share the common element of a serially connected user interface for low level configuration, control, and monitoring of the controller. Testing was performed with all controllers serially connected to a monitor so that all features provided through this interface could be evaluated.

4.2.2.1 Storage Concepts, Inc.: Concept 810

Storage Concepts provided a complete Concept 810 for testing. The Concept 810 is a 4+1, RAID 3 compatible controller and was provided as an example of how the Concept 910 (8+1 controller) might operate for an ORAID application. The unit tested included five Seagate 4 GB hard drives and the full complement of Storage Concepts features including Extended Differential Fast Bus (EDFB) port, Differential SCSI interface, and front panel display. The Concept 810 RAID controller board is comprised of a single, large (16-inch by 16.5-inch) printed circuit board. Drive connections are made from the top of the board to the disk drives. Disk drives are typically installed beneath the board and hard mounted to the chassis. Initial testing was performed using the installed hard drives.

The internal monitor was connected to determine the level of user configurability offered by the Concept 810. Although many configuration options are available, few were offered that directly support drive configuration. Once the unit was thoroughly tested with the magnetic drives, the optical drives were installed. Five Maxoptics drives were installed in place of the hard drives. The Concept 810 was able to recognize and spin-up all disk drives. Apparently, the ability to change drive configuration options is not required for the Concept 810 to recognize the optical drives.

The initial positive aspects of the Concept 810 quickly faded when testing began. Since the Sun and the PC both utilize single ended SCSI interfaces, a differential to single ended SCSI converter had to be obtained in order to connect the Concept 810 to the host platform. Once connected, the evaluation unit received would not communicate with either the SPARCstation or PC hosts. This problem persisted with the magnetic as well as optical disk drives. Support from Storage Concepts technical personnel proved of little help in bringing the unit on-line. This inability to communicate prevented any test data from being collected on the unit during the evaluation period. These problems combined with the lack of strong technical support from the vendor make it difficult to recommend the Concept 810 as the ORAID ADM controller. A presentation of the pros and cons of utilizing the Concept 810/910 are provided in Table 4-9.

Table 4-9. Analysis of Concept 810

PROS	CONS
Available 8+1 Configuration	Differential SCSI interface to host
Available EDFB Interface	External user interface is engineering oriented, not end-user oriented; difficult to determine disk parameters
	Evaluation unit is a 4+1, actual unit is a 8+1; not sure of how the actual unit will perform
	Physical design of evaluation unit and proposed unit will be difficult to integrate
	Requires extensive modification by manufacturer to support optical drives
	No front panel interface (keyboard) to system

4.4.4.2 Digi-Data Corporation: Z9000

The Digi-Data is a single, multi-module unit packaged as a standard, full height 5.25-inch disk drive. The Model Z9000 is a 4+1 controller supporting RAID levels 0, 3, and 5. The internal boards consist of a processor card, cache controller card, disk interface card, and front panel circuit card. For testing purposes, the Digi-Data was configured with 16 MB of RAM. Initial tests on the Digi-Data were performed using Seagate 4 GB hard drives. These drives were then exchanged for the Maxoptics optical disk drives. The Digi-Data controller was able to recognize the drives from the power-on state. Using the embedded Command Line Interface (CLI) feature, the controller was configured to match the operating performance of the optical drives. The Digi-Data CLI proved much more robust than the Storage Concepts internal monitor. Parameters such as transfer rate, transfer period, spin-up delay, and power-up delay could be adjusted through the CLI. Disk status including vendor information and capacity could also be directly viewed through the CLI. A presentation of the pros and cons of utilizing the Z900 are provided in Table 4-10

Table 4-10. Analysis of Z9000

PROS	CONS
External Command Line Interface (CLI) provides programmability and useful system information	CLI is only accessible via RS-232 port
Multiple RAID levels supported	Limited user feedback provided via front panel LCD
Physical size integrates easily into standard chassis	Maximum configuration for RAID 3: 4+1
Configurable to support reading and writing to optical drives	Requires manufacturers support to fully support optical features
Designed to be used as an OEM device; documentation and support are designed for the OEM	Limited front panel input (two buttons)

4.2.2.3 CMD Inc.: CRD-5000

The CMD CRD-5000 offered an excellent user interface, providing many programming features to the user through the front panel. However, the CRD-5000 consistently failed the configured Logical Unit (LUN) on spin-up. This LUN failure was determined to be related to the slow spin up time associated with the optical drives. The CRD-5000 does not allow the user to adjust the time out delay to compensate for slow disk drives. Thus, no meaningful test data could be generated using the CRD-5000. A presentation of the pros and cons of utilizing the CRD-5000 are provided in Table 4-11.

Table 4-11. Analysis of CRD-5000

PROS	CONS
Available as a 6+1 configuration	Would not communicate with optical drives without changes
Full menu-driven interface available through the front panel	Requires an ANSI or VT100 terminal to access interface externally
Excellent packaging	Little interest from OEM for changes without large volume
	Focus is on new product release targeted for Fall 95

4.2.3 Controller Evaluation and Selection

With all controller testing complete, a weighted comparison of controller characteristics was performed. The evaluation criteria were established based on factors that were important to the successful integration of a commercial RAID controller with optical disk drives. The criteria used to evaluate each controller are provided in Table 4-12.

Table 4-12. Description of Evaluation Criteria

Criteria	Description
Vendor Support of Modifications	Vendor support was considered critical for two primary reasons: initial controller testing and adapting the controller to the needs of optical disks for the final system
Number of OEM Modifications	Extensive modifications are expected to be costly and introduce unacceptable delays to the program development goals
Integration of Optical Drives	The level of support provided to the optical drives by the initial OEM configuration
Extent/Type of Modifications	The level of modifications was estimated to determine how much effort would be required by the vendor to meet ORAID requirements
Form Factor of Controllers	Physical configuration of the unit and its ability to be adapted to other mechanical configurations
User Interface	Amount of I/O provided to the user via the front panel

With the evaluation criteria established, a weight was assigned to each factor with respect to its importance in determining an overall score for each controller. The test results for each controller were then analyzed against each measurement criteria. Each controller was then assigned a value from 1 to 10 to indicate its relationship to the comparison criteria. A higher number indicates it met the expectations of the criteria in a favorable way. Table 4-13 illustrates the ranking criteria, weights and characteristic scores used to get a total comparative score for each controller. Based on these results, as well as extensive testing and analysis of each controller, the Digi-Data Z9000 is the clear winner. Therefore, this is the controller that will be used to implement the ORAID ADM.

Table 4-13. Controller Comparison Criteria

Factor	Weight	CMD	Digi-Data	Storage Concepts
Vendor Support	10	4	10	4
Number of Mods	8	6	8	7
Integration	7	5	9	3
Extent/Type of Mods	6	7	7	5
Form Factor	6	9	9	2
User Interface	4	8	8	2
Total Score		260	364	169

An analysis was then performed to determine what modifications would be required to fully support the optical drives through the Digi-Data Z9000 controller

4.2.3.1 Modifications to Improve Performance of Z9000

For an unmodified unit, the Z9000 performed remarkably well. However, since the unit is designed to support magnetic disk-based devices, there is significant room for improvement and grounds for concern in utilizing the Z9000. A preliminary list of these problems and concerns are provided in Table 4-14. Continued testing and analysis allowed the list to be further defined and categorized into areas for vendor modifications and areas for system support.

Table 4-14. Areas of Concern for COTS System Integration

Description of Concern/Problem	Resulting Limitation
Unit is optimized for performance with magnetic media	No support for optical media features; caching algorithms may not be optimized
No support for named media sets	No way of checking to determine if the installed disks are part of the same set
Maximum 4+1 configuration	Limits on-line access to 2.38 GBytes
Block transfer optimization for cache	Causes the transfer rate to drop for blocks greater than 4 MBytes

Vendor modifications were presented to Digi-Data for input on the best method to correct the potential deficiencies. This list is provided in Table 4-15.

Table 4-15. Digi-Data Z9000 Revisions

Description of Concern/Problem	Explanation
Control of cache control thresholds through mode page 0x02	An alternative method to adjust these values would also be very acceptable (i.e., CLI command, firmware mods, etc.). With write back cache enabled, once the cache is full, the Z-9000 will disconnect from the host for an extended period of time, usually causing the host to time-out. With magnetic drives, this does not seem to be a problem since the cache empties faster, thus the duration of the disconnect is less and time-out is avoided.
Reserved access (through the CLI pass-thru command) to one or two blocks on each disk in order to label media sets	Currently, the Z-9000 reserves 1024 sectors at the end of each disk for its own use. Since not all of these reserved sectors are used by the Z-9000 make use of one or two in order to accomplish the media set labeling task. Alternatively, more sectors could be reserved by the Z-9000 for this purpose and a media set labeling scheme be built into the Z-9000 firmware.
Implement an option in the pass-thru feature to send the same command to all disks (i.e., SCSI port number) on a single tier (i.e., SCSI ID), simultaneously	The capability to send the same command to all disks on all tiers, simultaneously to speed up system management
Firmware modifications to allow for the use of 1024 byte/sector media	Use of 1024 byte/sector media improves capacity and data transfer rate

As experienced with other RAID vendors, Digi-Data was reluctant to modify their firmware to meet all the needs to fully satisfy the ORAID ADM requirements. Therefore, work around features will be designed in as part of the system integration task.

4.2.4 OEM RAID Controllers: Lessons Learned

Manufacturers of RAID controllers were divided into two groups. Integrated system manufacturers who build complete systems based on their own proprietary RAID controller and manufacturers of OEM RAID controllers. When either group was contacted for information on their controller products with respect to optical drives, none were found who could validate performance with optical disk drives. Many would only offer guaranteed performance with a handful of magnetic disk vendors. Therefore, in order to determine any performance issues associated with using a magnetic based RAID controller in on optical system, unit testing of available products would have to be performed.

All controllers tested offer the end user the ability to set up and monitor the internal functions of the RAID controller. However, the interface to perform this function requires an external monitor. This is usually a dumb terminal connected to through a serial port to the controller. This interface tends to be awkward to use and understand. For many

users, locating the equipment to connect to the monitor port can be difficult. When questioned about the functionality of these monitor systems, most vendors replied that they were designed for the integrators to set up systems for end users, not as tools for end users to help control the system. This implied a built in dependency between the end user and the developer/integrator of the RAID system. It was therefore decided that ORAID would provide a better method to allow the end user to take advantage of built in monitor and control functions without relying on a third party for support.

Once all testing was completed, discussions were held with the respective controller vendors with respect to optical drive support. Most vendors saw no need to support optical disk technology. This lack of support was largely due to the lack of knowledge about the capabilities and growth in the optical storage industry. Once they were informed about the technology advances in the industry, their interest in supporting optical increased. However, they ultimately wanted to know what type of sales volume they could expect for any changes they made to their standard products. Without this information, they were unwilling to make any radical changes to their existing product. This position is typical for commercial product manufacturers who rely on profits from sales to drive engineering development.

4.3 System Design Considerations

With testing and component selection completed, the design focus shifted to creating a fully integrated system designed to meet the requirements of the ORAID ADM. Among the system design issues were storage, transfer rate performance, packaging, and integration with potential host systems. The system design concept was reviewed and a final integration path selected.


The principle focus of the ORAID development concentrated on fully implementing the ORAID with a commercial RAID controller with modification and design support provided by the OEM. Through the course of testing and analysis and with the selection of the Z9000, it became clear that this may not be the most practical approach to an ORAID implementation. Additional options needed to be considered before continuing with the development. The four options that were selected for further analysis are listed in Table 4-16.

Table 4-16. Design Approaches

Integration Approach	Description
Standard OEM Controller	Purchase and integrate a COTS, unmodified RAID controller to provide limited capabilities
Integrated COTS Components	Purchase separate COTS components and integrate utilizing developed software to smooth over integration issues
Standard OEM Controller with Embedded Control	Purchase a standard OEM RAID controller and integrate a COTS control module that can improve the performance of the RAID controller to more closely match ORAID requirements
Design and Develop Fully Compliant RAID Controller	Design a RAID controller capable of fully supporting the needs of optical drives and the performance goals of the ORAID

4.3.1 Evaluation of Design Options

The Digi-Data and the Maxoptics drives perform very well together. Preliminary tests proved that this basic system was able to perform marginally well as an Optical RAID. The results of this initial test are shown in Figure 4-15.

**Figure 4- 15. Preliminary Results of Integration Testing**

Although this initial performance graph appears to show basic system performance goals have been met, it fails to show other problems associated with this early integration effort. Such system level issues as support for features of optical media, media checking, and storage capacity must still be addressed. In considering other design options that might overcome some of the performance problems (Table 4-16, above), all options were

discussed in detail and the tradeoffs for each option carefully analyzed. A review and results of the analysis associated with each design option are shown in the following sections.

4.3.1.1 Option 1: Standard OEM Controller

This approach would end all further development and utilize an unmodified, COTS RAID controller. Severe system limitations, as discussed in prior sections, would result and the ORAID ADM would become a product with very few of the advantages the development was intended to provide. The resulting product would have little success in the commercial marketplace.

4.3.1.2 Option 2: Integrated COTS Components

An integrated COTS design would exclusively utilize COTS components to provide a work-around to some of the limitations provided by integrating a stock COTS RAID controller. The system would include multiple SCSI interface cards, a backplane, and a high level software executive to control the system.

Although technically feasible, this approach does have some developmental as well as operational risk. The host software used to manage multiple SCSI cards would be difficult to implement. In addition to host software concerns, the cards themselves would introduce problems. There are a limited number of IRQ and DMA channels available to the host. The more SCSI cards required, the more likely a conflict would occur. In addition, the SCSI cards perform on board processing of commands. This would make it difficult for the host to fully control the data transfer to member disks. The physical configuration of such a system could also be difficult.

4.3.1.3 Option 3: Integrated Support of an OEM RAID

Integrated support implies using external components to enhance the functionality of a COTS RAID controller. A combination of COTS and custom developed hardware and software would be implemented to work around the shortcomings of the selected controller. This method would make maximum use of the inherent features of the selected RAID controller by implementing an embedded control module to perform commands specific to optical disk drives. Although this approach is still limited by the features of the RAID controller, it is a fast, low-risk approach to meeting the goals of the ORAID ADM.

4.3.1.4 Option 4: Develop Custom RAID Controller

The high level of analysis performed in the development of the ORAID ADM has allowed Rising Edge to determine what features and functions a RAID controller must have in order to fully support optical disk drives. By evaluating commercial controllers, the best and worst features of commercial controllers have been determined. Rising Edge has also performed a full decomposition, analysis, and simulation of RAID Levels 0, 3, and 5.

These facets combined with previous hardware and software development experience, lead Rising Edge to believe this is the best way to achieve all the design goals for the ORAID ADM. However, this approach also has a marginal amount of technical risk not associated with the other approaches and would result in increasing the budget and extending the delivery date of the ORAID ADM.

Future generations of an Optical RAID would greatly benefit from the full development approach. Features that specifically support optical as well as magnetic features could be included. In addition, higher performance and capacity would be supported by increasing the number of drive channels. It is recommended that for long term success and acceptance of Optical RAID technology, that the custom developed approach be implemented in the future.

4.3.2 Mechanical Design Considerations

The mechanical design of the ORAID was driven by three main factors. It was determined during the analysis and testing of the Maxoptics optical disk drives that maintaining the air flow through the drive was critical for proper operation. Therefore the chassis design had to support a 3 cubic feet per minute airflow through the drives at all times. Secondly, in order to meet the ADM design goal of hot swappable disk drives, a device needed to be fabricated to allow the drives to be removed from the chassis. The final factor related to the cost of the chassis. It was hoped that a commercial, low cost chassis could be purchased to preclude the costs associated with a custom designed chassis.

4.3.2.1 Removable Drive Tray

With the large number of RAID systems on the market, it was anticipated that a removable drive tray could be purchased that would meet the needs of a 5.25-inch, full height disk drive. However, the majority of commercial trays are designed for half-height drives and do not allow access to the front of the installed disk drive. Front access is required in order to install and remove the optical media.

In order to overcome these limitations, it became necessary to design a removable drive tray. The goal of the tray design was to create a low-cost, functional unit that would allow the disk drives to be easily removed from the tray and replaced in the field. The tray needed to accommodate a printed circuit board that would hold connectors capable of supporting the hot swap function, allow access to the rear of the disk drive for cabling, and allow easy insertion and removal of the drive. The tray design is shown in Figure 4-16.

4.3.2.2 Chassis Design

Meeting the chassis requirements for ORAID also began with a review of commercially available components. Several vendors of RAID, PC and network server chassis were contacted. There were several common features discovered in the current crop of

commercially available chassis. Most RAID vendors are focused on supplying chassis and drive trays designed to accommodate 3.5-inch magnetic disk drives. Of the few vendors who supported full height 5.25-inch drives, most did not offer enough mounting space for the optical drives and controller or were designed to accommodate a specific manufacturers RAID controller. Standard PC chassis are generally too small and do not accommodate the required number of full height drives. This is largely true of the server chassis as well.




Figure 4- 16. Removable Drive Tray

Fortunately, a server chassis was discovered that met several basic criteria for the ORAID ADM chassis. The right side of the chassis is capable of accommodating five full-height

disk drives. The left side has sufficient space for the controller and the redundant power supplies. However, the front panel of the chassis still required modification. The decision to build or buy and modify was then reviewed.

In discussing the cost of chassis development with several sheet metal and plastic production shops, it became clear that it would require significant production volume to realize any cost savings with a newly developed chassis. It was determined that modifying a standard chassis would thus have significant cost advantages over a new design. The basic shell of the chassis could be maintained while the front panel could be replaced to support the required spacing of the new components. Internal components would then be added to provide support and structural integrity.

4.3.3 ORAID Power Supplies

Of the many commercial power supply vendors in the marketplace, only three were discovered who offered all the features desired for the ORAID power supply. The greatest limiting factor in qualifying supply vendors was in the physical size and configuration of the supply. The key features used in the final supply selection include sufficient power to support all system components, built-in load sharing for N+1 redundancy, designed for removability (i.e. hot-swap support), and status monitoring ports. Table 4-17 lists contacted vendors who had power supplies meeting these requirements.

Table 4-17. Power Supply Vendors

Vendor	Model
Omega	HP Series
SSI	PFQ 350
Sunpower Technologies	RPM Series

4.3.3.1 ORAID Power Budget

With the drives and RAID controller selected, a power budget was developed. This power budget was used to help select the ORAID system power supply. The calculated power budget is shown in Table 4-18.

In reviewing the power budget against the available supplies, it quickly became apparent that only one vendor could meet the ORAID ADM design goals. The Omega power supply was selected for use in the ORAID system. The HP1-2CD offers all the features required to support the needs of the ORAID ADM. The ORAID ADM design utilizes two supplies configured as fully redundant and supporting hot swap features. This design element will ensure a much higher level of availability for the ORAID than a single power supply system.

Table 4-18. Preliminary Power Budget

Power Supply Specification

Power Requirements:

Optical Drive:

	Tolerance	+5 VDC	+12 VDC
		+/- 5%	+/- 5%
Ripple		50 mV PP (max)	120 mV PP (max)
Current:	Typical	2.0 A	1.0 A
	Maximum	2.2 A	4.0 A

Controller:

	Tolerance	+5 VDC	+12 VDC
		+/- 5%	+/- 5%
Ripple		50 mV PP (max)	120 mV PP (max)
Current:	Typical	3.5 A	0.01 A
	Maximum	4.5 A	0.1 A

Total Load:

Source	Qty	+5 VDC	+12 VDC
Controller (typ)	1	3.5	0.01
Controller (max)	1	4.5	0.1
Optical Drive (typ)	5	2	1
Optical Drive (max)	5	2.2	4

5 Drive System

	+5 VDC	+12 VDC
typ	13.5	5.01
max	15.5	20.1
max Watts	77.5	241.2

Total: 319

4.3.4 System Design: Lessons Learned

There are always compromises associated with an integrated approach to system design. Although systems can be made to work, there are often improvements that could be performed to make the system better. When the integrated approach to system design is reviewed with respect to the program goals, time frame, and budget, the approach outline above serves to best meet contractual requirements.

In reviewing the method used to achieve the program goals two basic flaws in the system design approach were identified. First, it should have been realized early in the development cycle that a fully commercial solution could not be integrated by relying on

commercial vendor support. The RAID industry is so strongly focused on magnetic drive support and the revenues generated by such systems to be interested in alternatives. Second, the amount of "specmanship" associated with the published performance specifications for the optical disk drives created an unrealistic perception of system transfer performance. Drive performance, including magnetic hard drives, should be calculated based on transfer rate and media density or through real world testing. Vendor data sheets offer performance numbers that are theoretically achievable but do not directly translate to user applications.

4.4 Selection of an Integration Approach

In reviewing the possible options for creating a compliant ORAID ADM, the standard controller with embedded support option was chosen as the best approach for meeting current contract goals. This approach not only maintained the development schedule but allows improvements to the commercial RAID controller to be incorporated without relying on modifications from the controller manufacturer. As an example, the embedded support allows an intelligent user interface to be built that does not require external components for support. In order to create the most cost effective mix of commercial and developed elements, an evaluation of products capable of supporting this development goal was performed.

4.4.1 Embedded Controller Evaluation and Selection

In selecting an embedded controller for use in the system, several design criteria were established.

- The embedded controller would be selected from commercially available units. Little or no hardware design support would be required for successful integration of the unit.
- The embedded controller should have as small a footprint as possible to simplify integration into the final system.
- The embedded controller must conform to an industry standard. This will ensure a variety of peripheral options as well as a greater selection of manufacturers should a selected component become unavailable.
- Software development must be supported using readily available tools and development platforms. Special operating systems and development software are not desirable.
- The addition of the controller should have minimal impact on system cost.

Once the basic component level characteristics were determined, specific design requirements were established. Specifications for processing speed and power, memory, and input/output needs were determined.

Based on these criteria, components for the PC/104 bus standard and STD32 bus standard were evaluated. Of the two standards, the PC/104 offered more choices in type of processors available, number of vendors offering products, and price versus performance characteristics. Three vendors were chosen and their product line reviewed for components that would meet the development goals set for the embedded controller. A summary of the reviewed products is provided in Table 4-19.

Table 4-19. Evaluation of PC/104 Products

Source	Product
AMPRO	CoreModule /XTPlus 16 MHz, 16-bit NEC V51 CPU
Real Time Devices (RTD)	CMF8680 cpuModule 16 MHz, 16-bit C&T F8680 CPU
WinSystems	SAT V40 8 MHz, 16-bit NEC V40 CPU

After talking with sales representatives from each company, it was decided that the RTD card would best meet the needs of the ORAID ADM. The RTD card not only met all the preliminary design goals but also offered such features integrated LCD support, keypad scanning, and standard RAM configurations that met the anticipated needs of the developed code.

Code development for the RTD Module was divided into seven basic routines. The intent was to develop each module independently with appropriate interfaces to other modules or hardware components designed in place. The seven basic code modules are shown in Table 4-20.

Table 4-20. Embedded Code Modules

Routine	Description
RAID_MAIN	Main program and supporting utilities
PASSTHRU	SCSI pass-thru functions
DISPLAY	Sequence of display interactions and messages
COMM_LINE	Trapping and interpretation of Command Line Interface messages
ASYNC_IO	Message handling for all asynchronous input/output messages
ERROR	Error handling routines
MEDIA_HDL	Media handling routines

4.4.2 System Interface Component Design and Development

The initial goal of the embedded design approach was to purchase as many COTS components as possible. Although this was largely accomplished by selecting the embedded PC/104 module, there are still several hardware design elements required to complete the ORAID ADM. The two most system critical elements were defined as the User Interface Panel and the System Status and Monitoring Board. The User Interface Panel (UIP) is used to create the user interface. The panel consists of an LCD display and user entry keys. The UIP hardware design requires some low level interface circuitry, push-button switches, and an LCD display. The System Status Interface (SSI) uses a

register-based I/O and interrupt scheme to prompt an on-board embedded controller when an event has occurred within the RAID controller that requires the system to take some additional action. The interconnection of these components is illustrated in Figure 4-17.




Figure 4- 17. System Block Diagram

4.4.2.1 Completion of Developed Component Fabrication and Assembly

A complete list of circuit boards designed for use in the ORAID ADM and their associated status at the time of delivery is provided in Table 4-21. A brief description and summary of the status of each board is provided in the following sections.

Table 4-21. Developed Components

Part Number	Description	Status
S9002-01	Assembly, Disk Interface Backplane, disk tray	PCB, fabricated and tested
S9003-01	Assembly, Disk Interface Backplane, chassis	PCB, fabricated and tested
S9004-01	Assembly, Power Supply Distribution	PCB, fabricated and tested
S9005-01	Assembly, System Status Interface	PCB, fabricated and tested
S9006-01	Assembly, Display Push Button Switches	Hand fabricated and tested
S9007-01	Assembly, Display Control Interface	Hand fabricated and tested

4.4.2.2 System Status Interface (SSI)

The SSI provides the interface between the PC/104 and the SSI controller, the SSI and the User Interface Panel (UIP), and the SSI and the Z9000 RAID controller. The design currently utilizes the PC/104 for system control and as the display driver. The basic SSI design incorporates all the circuitry to eliminate the need for a PC/104 processor. Future revisions of the SSI will eliminate the PC/104 and rely solely on the as-built components to perform all system monitoring and control functions.

4.2.2.3 User Interface Panel

The User Interface Panel (UIP) is generically used to describe one of the two components comprising the System Information Display (SID). The UIP is used to provide an electrical interface between the LCD display, system push buttons, and the SSI. Because of the problems associated with contrast in the current LCD display, it was determined that a printed circuit design of the UIP would not be practical for the ORAID ADM. A new LCD will have to be selected to solve the contrast problem thus forcing a change to the UIP. A PCB for the UIP will be designed for future generations of the ORAID.

As a first attempt to improve the contrast, the software used to select display driving colors was modified. This technique provided an improved contrast ratio allowing the current system to function for the purposes of the ADM.

In order to improve future designs of the ORAID, the search for an alternate LCD was also completed. A strong candidate was found in the Stanley GMF32024BSLY, LED backlit display. This display is not only physically smaller, but also utilizes a higher refresh rate thus reducing the potential for a washed out appearance. In an effort to provide

better control of the LCD, the Epson 1330 LCD driver IC included on the SSI will be utilized in all future versions of the ORAID.

4.4.2.4 Power Supply Interface Board

The PSIB is a passive backplane that serves as a distribution point for all DC power signals from the load sharing power supplies. Critical design issues associated with the PSIB included internal power planes to handle the high current associated with the supplies, physical mating of the power supplies, and mechanical attachment to the chassis. These design challenges have all been met and the PSIB has been successfully integrated into the final ORAID ADM.

4.4.2.5 Disk Interface Boards

The Disk Interface Boards (DIBs) are a matched set of components designed to support the hot swap requirements for the optical disk drives. The most critical factor in designing the DIB is associated with the mechanical interface. The board positions must be carefully aligned so that a blind-mate installation can be supported. By utilizing individual backplane cards, small errors associated with mechanical alignment were eliminated.

4.4.3 Chassis Integration

The final chassis integration and assembly occurred shortly before delivery. The efforts to integrate and assemble this chassis became a much greater task than initially expected. Part of the problem stemmed from the lack of detail associated with the original mechanical design. Many support components that were required to build a single assembled unit needed to be designed and fabricated late in the program to facilitate integration and assembly. In addition, the chassis, as well as several of the received manufactured components, had to be manually modified to improve fit and function. Thus, the ability to consistently reproduce the ADM chassis is limited and will require a fair amount of redesign and review to create a full production version of the ORAID. A list of the basic mechanical components fabricated for the ORAID ADM is provided in Table 4-22.

In some instances, such as the ADM door (F9014-01), the delivered item could not be utilized the ORAID ADM due to design or manufacturing errors. In such cases, these components were not delivered as part of the final system. The necessity of such items will be reviewed and evaluated in future product development.

Table 4-22. Fabricated Parts, ORAID ADM Chassis

<i>Part Number</i>	<i>Description</i>
F9010-01	Fabrication, handle for drive tray
F9011-01	Fabrication, drive tray
F9012-01	Fabrication, chassis front panel
F9013-01	Fabrication, trim (top)
F9013-02	Fabrication, trim (bottom)
F9014-01	Fabrication, door
F9015-01	Fabrication, fill panel (drive)
F9015-02	Fabrication, fill panel (trim)
F9016-01	Fabrication, RAID controller chassis
F9017-01	Fabrication, RAID controller face plate
F9018-01	Fabrication, power supply frame
F9019-01	Fabrication, disk backplane frame
F9020-01	Fabrication, panel fill (left)
F9021-01	Fabrication, panel fill (right)
F9022-01	Fabrication, display housing
F9023-01	Fabrication, display panel

After receiving all mechanical components and attempting to assemble a complete ADM chassis, a review of the modified-COTS approach to chassis was conducted. The basic conclusion reached is that modified-COTS is acceptable for one-of-a-kind assembly, but the high level of rework required to make the components fit is unacceptable for production oriented approach to design. In order to move the ORAID ADM into a commercial sales marketplace, complete chassis redesign will have to be performed. Details associated with assembly problems are detailed in the following sections.

4.4.3.1 Modified Front Panel

A custom designed front panel was required for several reasons. The main reason for the redesign is to eliminate the vents in the existing panel to meet the airflow requirements of the disk drives. Secondly, the new panel allowed the hot swap features of the ORAID ADM to be supported. These include the optical disk drive, the power supplies, and the RAID controller. Although the front panel was designed to be assembled using pop rivets in existing mounting holes in the chassis frame, alignment of the front panel to the frame proved to be difficult.

When the chassis was assembled by the originally manufacturer, the front and rear panels were installed using a custom jig to hold the frame square during assembly. When the front panel is removed, the chassis loses stiffness and can flex. Without the assembly jig, installation of the new front panel relies on the tolerances of the machined parts and skill

of the assembler to recreate a truly square unit. The newly assembled ORAID ADM chassis is not truly square and thus causes binding and alignment problems with the hot swap components. This alignment must then be corrected in the assembly of the hot swap modules.

4.4.3.2 Removable Drive Trays

An early design goal of the drive trays was to create a removable module that would fit into a standard, unmodified, full-height disk drive bay. This stemmed from the desire to use COTS chassis components in the main ORAID ADM chassis. This design constraint limited the design options in terms of physical configuration and materials used.

The resulting drive tray does not have the structural integrity to support reliable hot swap use. The tray can flex horizontally and vertically. The handles and drive eject mechanism are difficult to use and do not add any structural integrity to the tray. The tray often binds when installed or removed from the chassis. In an effort to improve the installation and removability of the drive tray, guide pins were installed in the internal mounting frame to ensure any movement in the tray would not effect alignment of the circuit cards.

The existing drive tray will function for the ORAID ADM. However, these trays would not be suitable for a commercial unit. The drive trays would require redesign to improve stiffness and interface card alignment.

4.4.3.3 System Information Display

During the early system demonstration, the suggestion was made that a better method be considered for the LCD panel. In an effort to create a more integrated unit, the display was mounted into the top of the chassis. The display was installed at an angle to aid in visibility. A custom aluminum housing had to be utilized in order to meet the size and installation space constraints of the chassis.

4.4.4.4 RAID Controller Housing

The RAID Controller chassis was designed to house the commercial RAID controller and the custom designed SSI. The housing is a combination of new sheet metal and the existing commercial controller chassis. The combined unit requires three half-height drive openings in the chassis. The current design does not support hot swappability.

4.4.4.5 Mounting Frames

Special mounting frames were needed to support the hot swap power supplies and mounting for the Power Supply Interface Board (PSIB) as well the optical disk/Disk Interface Board (DIB). The power supply mounting frame was designed to allow the

supplies to slide easily into the frame and mate securely to the PSIB. The DIB mounting frame incorporates each individual backplane and guide pins to accommodate drive installation.

4.4.5 Lessons Learned

There are many advantages to utilizing commercial elements in any new design. The ORAID ADM is no exception. Selection of readily available components moved the design forward at a much faster pace than might have been realized by a full custom approach. However, application of the custom design also introduced many problems when the final integration phase began.

A better approach would be to use the COTS components that provide the best solution for their respective task and design all necessary components to achieve the desired integration results.

4.4.5.1 Integrated Electronics

The integration of the commercial PC/104 and the developed SSI proved to be a successful effort. However, this melding of the two technologies does not create an optimal production ready solution. The generic nature of the PC/104 card lacks several key elements that forced the design of the SSI to integrate a complete system. Primarily this is due to the intended application environment of the PC/104 architecture.

PC/104 cards are designed as modules which can be installed in a small enclosure and stacked to build a fully functional control computer. Therefore, each available card is designed to support a very specific, and somewhat specialized set of functions. In order to support the required functions of ORAID ADM, the PC/104 was found to be lacking in several areas. Either the stack height to achieve the desired functions was unacceptable or the needed support was not available.

The limitations of the PC/104 combined with the lack of certain features in the Z9000 largely determined the design criteria of the SSI. The SSI could have been a much better suited to meet the demands of the ORAID ADM if it had been designed with respect to system requirements rather than limitations and constraints of the existing architecture. It is believed that any future versions of the ORAID ADM will need to support a different style SSI that is a more integral part of the overall design.

Section 5 Analysis of Final Product

5.1 Comparison of Results Against Development Goals

5.1.1 Optical Disk Storage Subsystem

The final ORAID ADM successfully met all requirements for the selection of an optical disk storage device.

After successfully testing all available products, the Maxoptics T3-1300 standard commercial product offering was selected. The drive supported storage media consisting of a removable cartridge with the capacity to store a minimum of 650 MBytes per side for a total storage capacity of 1.3 Gigabytes. The drives are packaged in as a full-height configuration. The drive also supports the single-ended SCSI-2 data interface. Each drive was supplied with 4 MBytes of cache memory. Industry standard optical media (available from a number of vendors) was used with the drive.

5.1.2 RAID Controller

A commercial RAID controller was utilized in the ORAID ADM that was able to meet or exceed all of the program goals.

The Digi-Data Z9000 was implemented as the RAID engine in the ORAID ADM. The Z9000 proved flexible enough to support the timing parameters of the optical drives without modification. The Z9000 controller provided the following functions:

- SCSI-2 Fast and Wide (20 MBytes/second) interface between the host and the ORAID ADM
- support for RAID Levels 0, 3, and 5
- SCSI Fast (10 MBytes/second) interface between the disk and the RAID engine
- support for status monitoring through a DISKLED port

The ability to monitor system status and control the ORAID ADM utilizing a Z9000 feature known as the DISKLED port is supplemented by the custom designed circuitry of the SSI and UIP.

5.1.3 Equipment Chassis

The equipment chassis is a combination of commercial off the shelf components and custom designed mechanical assemblies. The chassis houses the optical disk drives, RAID controller, cooling fans and power supplies. Through the use of a custom designed drive tray, it is possible to remove or insert the optical drive unit without opening or otherwise altering the physical state of the chassis.

5.1.4 Host Control Platform

5.1.4.1 Host Processing Platform and Operating System

During the course of ORAID ADM development, several host processing platforms were used for testing. Primary testing was performed using PCs running Windows for Workgroups, Windows 95, Windows NT version 3.5.1 operating systems and workstation class computers running HP-UX and Sun Solaris 2.3 operating systems. All systems were capable of connecting and recognizing the ORAID ADM. Transfer performance was largely dictated by the SCSI interface card installed in the host platform. Depending on the operating system under test, single and multiple Gigabyte partitions were recognized. For final system delivery, the host workstation delivered is a Sun SPARCstation 5 running Sun Solaris 2.3.

5.1.4.2 Host Data Interface

Data transfer to and from the ORAID ADM is supported by a standard SCSI Fast and Wide interface. The ORAID ADM offers a host-independent interface, requiring no special drivers or software to connect to the host processing system.

5.1.5 Modes of Operation

The ORAID ADM supports two basic modes of operation. A brief description of each mode is provided in the following sections.

5.1.5.1 Dedicated to a workstation.

Dedicated operation is considered as operation of the system when the end user is working directly on the Workstation connected to the ORAID. A majority of the ADM testing was performed in this dedicated mode using a Sun SPARCstation 5 running Sun Solaris version 2.3. The ORAID should be backward compatible with other Sun operating systems.

5.1.5.2 File server on a client/server network.

File server operation is defined as a user environment where multiple workstations connect to the ORAID ADM host processing platform via a network interface to retrieve data files. Although no specific development was required to support this mode of operation, testing was performed to validate performance in a network environment. The ORAID ADM was tested using an interface office LAN without any difficulties.

5.1.5.3 System Software

The host processing platform connects to the ORAID ADM thorough a Fast and Wide SCSI-2 compatible interface. The operating system must support a standard SCSI-2

driver. The operating system is responsible performing all file management functions. The Solaris operating system was capable of recognizing the entire file volume.

5.1.6 Performance Goals

The system was designed against the performance goals shown in Table 5-1.

Table 5-1. Performance Goals and Results

Parameter	Goal	Results
Storage Capacity	6 Gigabytes	6 Gigabytes (3 GB per side)
Data Transfer Speed	50 Mbits/sec (burst)	65.5 Mbits/sec (read)
	25 Mbits/sec (sustained)	66.0 Mbits/sec (write) 22.69 Mbits/sec (read) 11.43 Mbits/sec (write)

Although the raw capacity meets the desired system goals, several other factors should be considered when comparing this number against system goals. The system uses five, 1.2 GByte pieces of media yet only half of the total capacity is available at any given time. The media must be manually removed and re-inserted into the ORAID to gain access to the entire media set. The application of RAID levels 3 or 5 further reduces the on-line capacity by 20% (i.e. one disk drive) to allow for parity storage. Media management information occupies 512 kBytes per disk. Finally, the operating system reduces the available storage space on the array.

For final testing, the ORAID ADM was connected to the SPARCstation and configured as a RAID 3 device. The 3.0 GB capacity was reduced to 2.4 GB for raw data storage. Media management data occupies 2.56 Mbytes of space and the operating system occupies an additional 17.4 MBytes. Therefore, the resulting available storage capacity is 2.38 GBytes.

Balanced read/write performance for the ORAID ADM became impossible to achieve due to the write performance of the optical disk drives. However, through the analysis phase of the development it was believed the sustained and burst transfer rates were expected to be met. This was not the case. Burst transfer rates were met by combining cache in the optical drives with cache in the RAID controller. Sustained transfer rates were measured over the entire usable storage space of the ORAID ADM. When tested this way the drives were unable to sustain the required transfer rate.

Several additional RAID system features were also included with the ORAID ADM. These include redundant, hot swappable power supplies and hot swappable disk drives.

5.1.7 Interface Specifications

The ORAID supports industry standard interfaces between the host processing platform and the ORAID ADM and within the ORAID ADM chassis. No custom or proprietary interfaces have been used. The host interface is a standard SCSI-2 Fast and Wide bus and can be purchased as either a single-ended or differential interface. A single ended interface was delivered with the ORAID ADM. The internal busses to the disk drives utilizes a standard single-ended SCSI Fast interface.

5.1.8 Maintenance

The ORAID ADM was designed to provide maintenance features that allow for easy removal and replacement of critical system components and meaningful diagnostic routines that display overall status of the system. The System Information Display provides control functions as well as monitoring of system components and diagnostics. Diagnostics displays include basic features such as media checks, disk status, and audible alarms for critical failures. The power supplies also support power good/power fail indicators as well as the ability to trigger the system alarm for an audible indication of failure.

5.2 Final Analysis

The ORAID ADM successfully proved the feasibility of grafting rewritable, magneto-optical disk drives within a RAID architecture. Over the course of this development, several issues became clear which must be considered in future technology efforts of this type.

The data storage industry is undergoing rapid advancement. Magnetic hard drives are revised and improved approximately every 6 months. Optical drives change approximately every 18 months. Future storage development projects need to consider where they might fall on this development cycle and how the design effort can best take advantage of the impending changes. The ORAID ADM successfully adopted an open architecture that allowed for an easy upgrade with respect to standard interfaces to the drives but was not open enough to adapt to the mechanical change of the optical drives.

The trade off between custom and COTS is a difficult decision. The ORAID ADM maximized the use of COTS across all components. In retrospect, this is probably not the best method for creating or integrating a new device. A better approach would be to allow COTS product to serve the system needs in the areas where they perform best. If they do not support the optimal performance goals of the system, then custom components should be designed that do offer optimal performance.

The optical disk drives are the prime example of a strong COTS product that could not be improved in the context of the design effort. Other components such as power supplies and the adoption of widely supported interface standards also proved valuable.

In the case of the ORAID ADM, this concept of COTS vs. optimal performance applies specifically to the OEM RAID controller used in the system. Although the Digi-Data Z9000 proved to be the best available device to meet the goals of the development, it lacked many features and functions that could have led to a better ORAID device. The work around for its limitations resulted in additional COTS boards as well as a custom designed circuit to perform all system functions. The result was a less than desirable, yet functional, implementation of the system.

If the design goals had been refocused to eliminate the strong desire for COTS in favor of optimal performance, a better final product would have been produced that more globally met the goals of the contract. The risk associated with the custom development of the controller can be considered small in terms of the compromises that were made in system functionality and performance.

Overall, the ORAID ADM can be viewed as proving the concept of optical disk based RAID storage is viable. Future improvements in technology should allow the system to fully meet the performance goals of the contract.

Section 6 Redesign and Refinement

6.1 Changes in the Optical Drive Market

As the final delivery of the ORAID ADM was being completed, a tremendous shift occurred in the optical disk drive market. All major drive manufacturers announced a new versions of their product offering enhanced features, double the storage, half the size, and double the transfer rate. In an effort to reanalyze the ORAID ADM in light of these changes and improve on the development efforts associated with the first article, a second unit was developed.

The extensive amount of research, analysis, and test performed for the first article ADM proved invaluable in reducing the design and delivery time for the second generation ORAID ADM. It was believed that by using the first generation architecture as a baseline and taking advantage of the new drive products, the second generation ORAID ADM would achieve the intended system design goals. With this in mind, redesign was performed to incorporate the new technology, take advantage of lessons learned during the development of the first article, and improve on the overall system.

6.1.1 Selecting an Optical Drive

Three classes of product improvement were announced. Industry wide, the basic 5.25-inch MO drive doubled in capacity (2.6/2.3 GBytes), reduced in height (standard half-height device), and improved transfer speeds (doubling of read/write performance). This improved drive is available from most of the vendors previously contacted during the original development. Nikon Precision, Inc. took this new drive and added an additional feature labeled Light Intensity Modulation Direct Overwrite (LIM DOW). LIM DOW allows optical drives to write data in one pass, thus theoretically offering balanced read/write performance. The final product increased capacity by almost 80% over the existing products. Pinnacle Micro announced their Apex Optical Hard Drive offering 4.6 GB of capacity. Drives from each of these categories were requested for evaluation.

6.1.2 Testing

Drive testing was conducted primarily to verify the availability and performance claims of a select group of manufacturers. The detailed testing performed under the initial contract development was not conducted. In general, the new 2.6 GB optical drives met the expectations of the second generation ORAID ADM. However, the desired goal was to implement either the higher performance LIM DOW drives or the higher capacity Apex drives in the final deliverable.

Our previous relationship with Pinnacle allowed us to get early versions of the Apex for evaluation. The Apex uses non-standard media. Pinnacle Micro, through Mitsubishi Chemical, are the only company manufacturing this high capacity media. This media will not work in any other MO drive. In addition, the Apex is only backward compatible to the 2.6/2.3 GB media. It will not work with early generations of optical media. Although this is not a real concern with ORAID, since media must always be used as a full set in the device, it may effect the overall marketability of the Apex.

It was clear from these units that Pinnacle announced the product well ahead of its ability to meet the engineering challenges associated with the drive. Early units failed repeatedly and those that did work were far from their performance specifications. It was determined through this process that these drives would not be ready in time to meet system delivery deadlines.

Nikon was contacted to obtain evaluation units of their LIM DOW drive. The Nikon supports standard, non-LIM DOW 2.6/2.3 GB media as well as older generation of MO media. As with the Apex, the drive was far from mature when tested. Several issues were uncovered with their firmware that caused poor write performance and SCSI compatibility problems with the drive. By working with Nikon, we were able to help resolve several firmware issues. However, the write performance never achieved more than a 50% increase over standard 2.6 GB drives. Therefore, this drive was also eliminated as a candidate.

The Maxoptics T4-2600 was selected as the storage device for the second generation ORAID ADM. The performance, size, and capacity of the drive coupled with previous experience with Maxoptics and their T3-1300 drive made this a natural choice. The T4-2600 is a 5.25-inch, half-height optical storage device. It accepts removable, rewritable optical disk cartridges in five capacities: 2.6 GB, 2.0 GB, 1.3 GB, 1.0 GB (read only) and 650 MB. The drive supports a standard SCSI-2 interface. Although currently not available with LIM DOW technology, Maxoptics plans to release a T5 version of the drive during 1997 that will offer this feature.

6.2 Adapting Lessons Learned

6.2.1 Optical Disk Drives

Drive testing was considered the only acceptable method for determining system performance. Previous experience indicated that the published specifications do not reflect the actual data transfer performance of the drives. Therefore, only those drives described above were selected for testing. These drives were primarily selected based on experience with previous versions, availability, and vendor support. These three companies also represent the leaders in 5.25-inch optical technology, often branding their products for other companies such as Plasmon, Most, and Olympus.

6.2.2 OEM RAID Controller

The Digi-Data Z9000 RAID controller was also used in second generation ORAID ADM. Although the Z9000 is scheduled to be discontinued in late 1997, previous testing indicates that other multi-channel RAID controllers are currently not available to meet the demands of optical disk drives. Although several new controllers have been introduced to the market since the initial test were performed, most of them do not support enough channels or the desired RAID level required by the ORAID ADM.

Rather than test a new array of controllers, it was decided that the Z9000 would be used. Although future generations of the ORAID will require investigation and selection of a new controller, the current solution will provide excellent performance for the near term.

6.2.3 System Design

The overall system design represents the largest single change of scope to the ORAID ADM design. Electrically, the design remains largely unchanged. The application of an industry standard SCSI architecture early in the design allowed an easy migration path to the new optical disk drives. The RAID controller was not changed and therefore little redesign effort was focused on the this part of the system.

The system packaging became the major focus of the redesign effort. It was believed that a better production-focused chassis could be developed from a ground up approach. The requirements for the chassis were established to support accessibility and removability of components, and expansion.

The chassis design was formed around an integral card cage that would support the hot swap requirements of the system components as well as the physical dimensions of the optical disk drives. Once this card cage was defined, the exterior of the chassis was designed and all necessary brackets and component mounting pieces were completed. A complete list of designed parts is provided in Table 6-1.

Table 6-1. Fabricated Chassis Components

Part Number	Description	Status
PL0001010-01	Guide, Left	Complete
PL0001011-01	Cover, Left/Right	Complete
PL0001012-01	Bracket, Support, Upper	Complete
PL0001013-01	Bracket, Fan	Complete
PL0001014-01	Plate, Cosmetic, Left/Right	Complete
PL0001015-01	Plate, Cosmetic, Center	Complete
PL0001016-01	Plate, Cosmetic, Lower	Complete
PL0001017-01	Bracket, Support, Lower	Complete
PL0001018-01	Bracket, Power Interface	Complete
PL0001019-01	Front Panel, Optical Drive	Complete

Part Number	Description	Status
PL0001023-01	Cover, Power Supply, Top	Complete
PL0001026-02	Block, Controller, Left/Right Side	Complete
PL0001027-01	Plate, Drive, Lower	Complete
PL0001027-02	Cover, Optical Drive	Complete

The lessons learned from the original design proved valuable in creating a better package for the second ORAID ADM. Unfortunately, the final package is less than desirable when viewed as a long term solution for production. Among the chassis faults are a high cost to reproduce, difficult and lengthy assembly process, excessive weight that would ultimately effect shipping costs, and limited room for external I/O cabling.

Although the chassis offers some improvements over the modified COTS server chassis, a better approach to chassis design would be to model the design style of the commercial chassis and add the special requirements of the ORAID ADM. Commercial chassis are designed for low cost production and fast assembly. Although they achieve some of their cost advantage through high quantity, there are certain traits common to all chassis that greatly reduce costs. Examples include:

- Access to interior components through easily removable panels. In many cases, a one piece skin or pull off door is used that is totally attached from the outside of the chassis.
- Simple, riveted internal card cage punched to accept standard component sizes.
- I/O panels attached to openings in the back of the chassis to allow for customization/modification of units.

Based on the as delivered chassis, additional rework will be required to create a mass market style chassis. Many components can be reused, but redesign is still the best option.

The chassis was designed to support several electronics parts and assemblies used in the ORAID ADM. These assemblies combine circuit cards and cables to make a complete functional module. A list of ORAID ADM assemblies is provided in Table 6-2.

Table 6-2. List of ORAID ADM Assemblies

Part Number	Description	Status
PL0023001-01	PL, Assembly, Desk Side Chassis	Complete
PL0023002-01	PL, Assembly, RAID Controller Module	Combined components
PL0023003-01	PL, Assembly, Power Supply Module	Hand built load share PCB
PL0023004-01	PL, Assembly, Disk Carrier Module	Complete with PCB fab
PL0023005-01	PL, Assembly, System Info Display Module	PCB and hand built parts
PL0023006-01	PL, Assembly, DC Fan Tray	Complete

Part Number	Description	Status
PL0023007-01	PL, Assembly, RAID Controller Housing	Complete
PL0023101-01	PL, CCA, Data Distribution Backplane	Complete
PL0023102-01	PL, CCA, Disk Interface Backplane	Complete
PL0023103-01	PL, CCA, RAID Controller Backplane	Schematic and PL
PL0023104-01	PL, CCA, Keypad and Display Interface	Hand built
PL0023106-01	PL, CCA, Control Backplane	Schematic and PL
PL0023107-01	PL, Assembly, Overlay, Front Panel	Schematic and PL
PL0023111-01	PL, Assy, Cable, User Interface/Backplane	Complete
PL0023112-01	PL, Assy, Cable, User Interface/IDC	Complete
PL0023201-01	PL, PWA, Data to Disk Interface Backplane	Complete

In an effort to complete the system delivery on time, several hand-built components were delivered in the final ORAID ADM. In addition, mechanical design errors in the connector pin assignment on the Disk Interface Backplane had to be corrected with cables between the Backplane and the RAID controller. Ultimately, these design errors will be corrected when the Control Backplane is fabricated.

6.2.4 System Interface

The system interface remained largely unchanged. The LCD display was upgraded to the Stanley LED back-lit unit discovered during the original ORAID ADM design. In order to support this display, the original SSI and UIP were replaced with the Rising Edge Technologies Intelligent Display Controller (IDC).

The IDC was developed by Rising Edge Technologies as a standard product to support I/O processing, interface, and display driving for a wide range of control system applications. This product, integrated into the ORAID ADM, allows for a better user interface, improved error handling and response, reduced cost, and better reliability. Several display options are available using this controller as well as the ability to design a custom GUI through reprogramming. These options are not possible with the multi-card PC/104 approach used on the first ORAID ADM.

6.3 Analysis of Final Product

The second generation ORAID ADM was designed to retain all of the key user interface, connectivity, and maintenance characteristics of the original system. When compared to the original performance goals, the second generation ORAID ADM moved significantly towards meeting the contract expectations. Table 6-3 lists the new performance numbers.

Table 6-3. Performance Goals and Results

Parameter	Goal	Results
Storage Capacity	6 Gigabytes	11.5 Gigabytes (5.75 GB per side)
Data Transfer Speed	50 Mbits/sec (burst)	94.5 Mbits/sec (read) 47.5 Mbits/sec (write)
	25 Mbits/sec (sustained)	47.5 Mbits/sec (read) 24 Mbits/sec (write)

It is interesting to note that the new system can improve in the capacity and performance numbers by substituting in an alternate optical hard drive. For example, using the Pinnacle Micro Apex drive, system capacity can be increased to 21.5 GB (10.75 GB per side) while the Nikon LIM DOW drive can double the write performance. Although these drives did not perform well during the initial test phase for this system, production models shipped in early 1997 largely met the original specifications of the drive manufacturers. It is possible that future systems might be purchased to take advantage of the benefits these drives bring to the system.

The benefits and performance available through the implementation of an Optical RAID storage system remain strong. The second deliverable moved closer to establishing this technology as a mainstream storage device. Early feedback from customers in the document imaging, medical records management, and video production industries support the viability of this product.

The ORAID ADM has established some basic features and functions that will help ensure its success. The key factor is the implementation of a widely used interface such as the SCSI bus. This ensures an easy path to adopt changes in optical storage technology. The mechanical change to support half-height drives allows the system to be reconfigured in a commercial tower chassis as long as the hot swap drive trays are not used. This increases the potential market for the product by reducing cost for end-users who may not need 24 hour-a-day/7 day-a-week data protection afforded by hot-swap components.

Some additional refinement is required to make ORAID a stronger commercial product. As stated early, the majority of this refinement is associated with the packaging. However, based on the performance of the product and the extensive amount of positive customer feedback, Optical RAID appears to be a viable candidate for a piece of the commercial data storage product marketplace.

6.4 Further Development and Enhancements

The ORAID ADM accomplished several positive steps in proving the application of optical storage in a parallel access array such as RAID. As a barometer for the market, ORAID continues to generate a positive response. Numerous applications and different operating scenarios have been suggested for this technology. It offers performance,

archive, and reliability features previously unavailable in the optical storage market. However, there are still a large number of improvements and enhancements that could be made to make the system more marketable, increase its performance and capacity, and provide a future path for component changes and standards adaptation.

The ORAID ADM and the subsequent market exposure it has received point out two areas that should be pursued to improve the product design. The first is related to the RAID control engine and the second is application driven.

The RAID controller engine is a hybrid development created from a two COTS products designed to work together. Although this approach is effective and functional, it limits features that are desirable in a removable media system. The limitations are a direct result of the Digi-Data Z9000 controller. This controller is optimized to support magnetic hard disk drives. Through a combination of specific firmware changes to the Z9000 and the integration of a Rising Edge control board, the basic functionality of an Optical RAID has been obtained. Further research is required to determine what features, functions, and performance enhancements a RAID controller designed with optical hard drives in mind could afford such a system. It is our belief there a significant number of desirable features which could be obtained. A few areas that would be improved include:

- Improved cold/warm boot performance
- Improved error handling
- Improved media control
- Increased capacity and better data transfer performance

In terms of application support, ORAID offers the ability to provide high performance near-line storage to applications requiring large (multi-GByte) data access. By teaming ORAID with additional software and hardware components, and improve data management system might be created that would greatly benefit such applications.

Future efforts focused on improving ORAID should review and define methods to achieve these, and other, performance goals.

Appendix: Optical Drive Test Results

Optical Drive Description:

Olympus MOS525E Standard, 5.25 inch, half height

Interface: Single Ended Cache: 1 MB Rotational Speed: 3500 RPM

Power Requirements:

	+5 VDC	+12 VDC
Tolerance	+/- 0.25 V	+/- 0.6 V
Ripple	100 mV PP (max)	100 mV PP (max)
Current: Typical	1.2 A	1.0 A
Maximum	1.4 A	2.1 A

2 Pass Write: Erase/Write

3 Pass Write: Erase/Write/Verify

Media Description:

Maxell MA-132-S1 Format: ZCAV Type: 512 Byte/sector Formatted Capacity: 1.19 GB

Read Test:

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Transfer Rate
4544	128000	5.82E+08	1110.49
9088	64000	5.82E+08	847.30
18176	32000	5.82E+08	618.43
36531	16000	5.84E+08	429.42
72702	8000	5.82E+08	301.86
145404	4000	5.82E+08	223.31
290808	2000	5.82E+08	153.99

Write Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode	X			
Write Cache				X
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Transfer Rate
4544	128	581632	329.70
9088	64	581632	310.37
18176	32	581632	232.51
36531	16	584496	141.55
72702	8	581616	83.93
145404	4	581616	47.10
290808	2	581616	28.16

Write Test 2:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache				X
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Transfer Rate
4544	128	581632	491.87
9088	64	581632	451.94
18176	32	581632	326.17
36531	16	584496	204.83
72702	8	581616	112.33
145404	4	581616	59.50
290808	2	581616	37.83

Write Test 3:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode	X			
Write Cache			X	
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Cache Disabled	Cache Enabled
4	131072	524288	221.65	664.94
8	131072	1048576	266.67	395.37
12	131072	1572864	266.20	417.39
8	65536	524288	239.25	445.22
16	65536	1048576	235.94	344.78
24	65536	1572864	233.08	397.67
16	32768	524288	190.33	282.87
32	32768	1048576	190.33	238.69
48	32768	1572864	191.52	228.91
32	16384	524288	137.27	163.58
64	16384	1048576	134.21	162.28
96	16384	1572864	136.41	164.45
64	8192	524288	73.46	104.60
128	8192	1048576	73.67	100.79
192	8192	1572864	74.96	100.26
128	4096	524288	45.03	57.85
256	4096	1048576	44.81	55.50
384	4096	1572864	44.82	56.72
256	2048	524288	24.08	40.51
512	2048	1048576	24.08	38.12
768	2048	1572864	23.96	38.00

Write Test 4:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache			X	
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Cache Disabled	Cache Enabled
4	131072	524288	2.91E+05	6.65E+05
8	131072	1048576	3.97E+05	7.16E+05
12	131072	1572864	3.94E+05	5.09E+05
8	65536	524288	3.46E+05	3.88E+05
16	65536	1048576	3.52E+05	4.15E+05
24	65536	1572864	3.49E+05	3.67E+05
16	32768	524288	2.74E+05	3.22E+05
32	32768	1048576	2.78E+05	3.27E+05
48	32768	1572864	2.74E+05	3.21E+05
32	16384	524288	1.98E+05	2.23E+05
64	16384	1048576	2.03E+05	2.07E+05
96	16384	1572864	2.03E+05	2.15E+05
64	8192	524288	1.11E+05	1.43E+05
128	8192	1048576	1.12E+05	1.34E+05
192	8192	1572864	1.11E+05	1.50E+05
128	4096	524288	5.52E+04	7.06E+04
256	4096	1048576	5.53E+04	7.04E+04
384	4096	1572864	5.54E+04	7.15E+04
256	2048	524288	3.01E+04	5.21E+04
512	2048	1048576	3.00E+04	4.99E+04
768	2048	1572864	2.98E+04	5.05E+04

Band Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache				X
SCSI Parity				X

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		194		131072		2.54E+07		4.02E+05
1		206		131072		2.70E+07		4.23E+05
2		218		131072		2.86E+07		4.33E+05
3		231		131072		3.03E+07		4.67E+05
4		243		131072		3.19E+07		4.77E+05
5		255		131072		3.34E+07		4.86E+05
6		267		131072		3.50E+07		4.85E+05
7		279		131072		3.66E+07		5.45E+05
8		291		131072		3.81E+07		5.58E+05
9		303		131072		3.97E+07		5.54E+05
10		315		131072		4.13E+07		5.57E+05
11		327		131072		4.29E+07		5.82E+05
12		340		131072		4.46E+07		4.71E+05
13		352		131072		4.61E+07		4.70E+05
14		364		131072		4.77E+07		4.73E+05
15		367		131072		4.81E+07		4.81E+05

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		388		65536		2.54E+07		3.52E+05
1		412		65536		2.70E+07		3.57E+05
2		436		65536		2.86E+07		3.68E+05
3		461		65536		3.02E+07		4.09E+05
4		485		65536		3.18E+07		4.11E+05
5		509		65536		3.34E+07		4.29E+05
6		533		65536		3.49E+07		4.33E+05
7		558		65536		3.66E+07		4.40E+05
8		582		65536		3.81E+07		4.38E+05
9		606		65536		3.97E+07		4.47E+05
10		630		65536		4.13E+07		4.56E+05
11		654		65536		4.29E+07		5.06E+05
12		679		65536		4.45E+07		5.16E+05
13		703		65536		4.61E+07		5.24E+05
14		727		65536		4.76E+07		5.47E+05
15		733		65536		4.80E+07		5.40E+05

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	775	32768	2.54E+07	2.76E+05
1	824	32768	2.70E+07	2.79E+05
2	872	32768	2.86E+07	2.83E+05
3	921	32768	3.02E+07	2.79E+05
4	969	32768	3.18E+07	2.89E+05
5	1018	32768	3.34E+07	3.07E+05
6	1066	32768	3.49E+07	2.96E+05
7	1115	32768	3.65E+07	2.99E+05
8	1163	32768	3.81E+07	2.96E+05
9	1211	32768	3.97E+07	3.18E+05
10	1260	32768	4.13E+07	3.76E+05
11	1308	32768	4.29E+07	3.64E+05
12	1357	32768	4.45E+07	3.72E+05
13	1405	32768	4.60E+07	3.72E+05
14	1454	32768	4.76E+07	3.87E+05
15	1465	32768	4.80E+07	3.97E+05

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	1550	16384	2.54E+07	2.03E+05
1	1647	16384	2.70E+07	2.04E+05
2	1744	16384	2.86E+07	1.94E+05
3	1841	16384	3.02E+07	1.91E+05
4	1938	16384	3.18E+07	1.92E+05
5	2035	16384	3.33E+07	1.97E+05
6	2132	16384	3.49E+07	2.09E+05
7	2229	16384	3.65E+07	1.97E+05
8	2325	16384	3.81E+07	2.11E+05
9	2422	16384	3.97E+07	2.12E+05
10	2519	16384	4.13E+07	2.01E+05
11	2616	16384	4.29E+07	2.06E+05
12	2713	16384	4.44E+07	2.07E+05
13	2810	16384	4.60E+07	2.06E+05
14	2907	16384	4.76E+07	2.08E+05
15	2930	16384	4.80E+07	2.22E+05

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	3100	8192	2.54E+07	1.11E+05
1	3294	8192	2.70E+07	1.07E+05
2	3488	8192	2.86E+07	1.13E+05
3	3682	8192	3.02E+07	1.13E+05
4	3875	8192	3.17E+07	1.10E+05
5	4096	8192	3.36E+07	1.14E+05
6	4263	8192	3.49E+07	1.15E+05
7	4457	8192	3.65E+07	1.13E+05
8	4650	8192	3.81E+07	1.10E+05
9	4844	8192	3.97E+07	1.15E+05
10	5038	8192	4.13E+07	1.17E+05
11	5232	8192	4.29E+07	1.17E+05
12	5425	8192	4.44E+07	1.14E+05
13	5619	8192	4.60E+07	1.11E+05
14	5813	8192	4.76E+07	1.10E+05
15	5859	8192	4.80E+07	1.10E+05


Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	6200	4096	2.54E+07	5.52E+04
1	6588	4096	2.70E+07	6.07E+04
2	6975	4096	2.86E+07	5.90E+04
3	7363	4096	3.02E+07	5.75E+04
4	7750	4096	3.17E+07	5.67E+04
5	8138	4096	3.33E+07	6.01E+04
6	8525	4096	3.49E+07	5.93E+04
7	8913	4096	3.65E+07	6.01E+04
8	9300	4096	3.81E+07	6.27E+04
9	9688	4096	3.97E+07	5.97E+04
10	10075	4096	4.13E+07	5.98E+04
11	10463	4096	4.29E+07	6.04E+04
12	10850	4096	4.44E+07	6.01E+04
13	11238	4096	4.60E+07	5.84E+04
14	11625	4096	4.76E+07	5.78E+04
15	11717	4096	4.80E+07	6.25E+04


Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		12400		2048		2.54E+07		3.01E+04
1		13175		2048		2.70E+07		2.94E+04
2		13950		2048		2.86E+07		3.02E+04
3		14725		2048		3.02E+07		3.69E+04
4		15500		2048		3.17E+07		3.40E+04
5		16275		2048		3.33E+07		3.37E+04
6		17050		2048		3.49E+07		3.58E+04
7		17825		2048		3.65E+07		4.17E+04
8		18600		2048		3.81E+07		4.17E+04
9		19375		2048		3.97E+07		4.32E+04
10		20150		2048		4.13E+07		4.14E+04
11		20925		2048		4.29E+07		4.12E+04
12		21700		2048		4.44E+07		4.12E+04
13		22475		2048		4.60E+07		4.10E+04
14		23250		2048		4.76E+07		4.04E+04
15		23433		2048		4.80E+07		4.01E+04


Olympus MOS525E Average Read Rate




Olympus MOS525E Average Write Data Rates




Olympus MOS525E 128k Block 3 Pass Write Cache Test




Olympus MOS525E 64k Block 3 Pass Write Cache Test




Olympus MOS525E 32k Block 3 Pass Write Cache Test




Olympus MOS525E 16k Block 3 Pass Write Cache Test




Olympus MOS525E 8k Block 3 Pass Write Cache Test



Olympus MOS525E 4k Block 3 Pass Write Cache Test




Olympus MOS525E Band Testing (2-Pass Write)




Olympus MOS525E 128k Block 2 Pass Write Cache Test

■ Cache Disabled ■ Cache Enabled




Olympus MOS525E 64k Block 2 Pass Write Cache Test


■ Cache Disabled ■ Cache Enabled




Olympus MOS525E 32k Block 2 Pass Write Cache Test




Olympus MOS525E 16k Block 2 Pass Write Cache Test




Olympus MOS525E 8k Block 2 Pass Write Cache Test




Olympus MOS525E 4k Block 2 Pass Write Cache Test



Olympus MOS525E 2k Block 3 Pass Write Cache Test



Olympus MOS525E 2k Block 2 Pass Write Cache Test



Optical Drive Description:

Hewlett-Packard C1716T Standard, 5.25 inch, full height

Interface: Single Ended Cache: 512k Rotational Speed: 2400 RPM

Power Requirements:

	+5 VDC	+12 VDC
Tolerance	+/- 5%	+/- 5%
Ripple	100 mV PP (max)	100 mV PP (max)
Current: Typical	2.25 A	1.0 A
Maximum	2.5 A	3.0 A

2 Pass Write: Erase/Write

3 Pass Write: Erase/Write/Verify

Media Description:

Maxell MA-132-S1 Format: ZCAV Type: 512 Byte/sector Formatted Capacity: 1.19 GB

Read Test:

Read Cache Enabled

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	729.45
9088	64	581632	630.85
18176	32	581632	516.51
36531	16	584496	346.94
72702	8	581616	238.99
145404	4	581616	141.43
290808	2	581616	123.47

Read Cache Disabled

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	711.51
9088	64	581632	617.50
18176	32	581632	502.01
36531	16	584496	347.53
72702	8	581616	240.41
145404	4	581616	138.84
290808	2	581616	74.64

Write Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache			X	
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	410.48
9088	64	581632	464.12
18176	32	581632	502.59
36531	16	584496	440.48
72702	8	581616	370.35
145404	4	581616	261.43
290808	2	581616	173.64

Write Test 2:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache				X
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	351.39
9088	64	581632	296.58
18176	32	581632	241.86
36531	16	584496	170.90
72702	8	581616	133.30
145404	4	581616	71.93
290808	2	581616	37.89

Conditions	3 Pass	2 Pass	Read	Write
Write Mode		X		
Test Mode				X
Burst Size: unlimited				X

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Cache Disabled	Cache Enabled
4	131072	524288	2.75E+05	6.65E+05
8	131072	1048576	2.67E+05	3.95E+05
12	131072	1572864	2.69E+05	4.17E+05
8	65536	524288	2.46E+05	4.45E+05
16	65536	1048576	2.45E+05	3.45E+05
24	65536	1572864	2.45E+05	3.98E+05
16	32768	524288	1.94E+05	2.83E+05
32	32768	1048576	1.90E+05	2.39E+05
48	32768	1572864	1.89E+05	2.29E+05
32	16384	524288	1.37E+05	1.64E+05
64	16384	1048576	1.35E+05	1.62E+05
96	16384	1572864	1.36E+05	1.64E+05
64	8192	524288	1.22E+05	1.05E+05
128	8192	1048576	1.30E+05	1.01E+05
192	8192	1572864	1.30E+05	1.00E+05
128	4096	524288	7.17E+04	5.79E+04
256	4096	1048576	7.22E+04	5.55E+04
384	4096	1572864	7.20E+04	5.67E+04
256	2048	524288	3.59E+04	4.05E+04
512	2048	1048576	3.57E+04	3.81E+04
768	2048	1572864	3.56E+04	3.80E+04

Write Cache Test 1:

Conditions	3 Pass	2 Pass	Read	Write
Write Mode		X		
Test Mode				X
Burst Size: 16384				X

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Cache Disabled	Cache Enabled
4	131072	524288	2.75E+05	6.65E+05
8	131072	1048576	2.67E+05	7.16E+05
12	131072	1572864	2.69E+05	5.09E+05
8	65536	524288	2.46E+05	3.88E+05
16	65536	1048576	2.45E+05	4.15E+05
24	65536	1572864	2.45E+05	3.67E+05
16	32768	524288	1.94E+05	3.22E+05
32	32768	1048576	1.90E+05	3.27E+05
48	32768	1572864	1.89E+05	3.21E+05
32	16384	524288	1.37E+05	2.23E+05
64	16384	1048576	1.35E+05	2.07E+05
96	16384	1572864	1.36E+05	2.15E+05
64	8192	524288	1.22E+05	1.43E+05
128	8192	1048576	1.30E+05	1.34E+05
192	8192	1572864	1.30E+05	1.50E+05
128	4096	524288	7.17E+04	7.06E+04
256	4096	1048576	7.22E+04	7.04E+04
384	4096	1572864	7.20E+04	7.15E+04
256	2048	524288	3.59E+04	5.21E+04
512	2048	1048576	3.57E+04	4.99E+04
768	2048	1572864	3.56E+04	5.05E+04

Block Write Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache				X

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		194		131072		2.54E+07		2.73E+05
1		206		131072		2.70E+07		2.76E+05
2		218		131072		2.86E+07		2.83E+05
3		231		131072		3.03E+07		2.83E+05
4		243		131072		3.19E+07		2.87E+05
5		255		131072		3.34E+07		3.10E+05
6		267		131072		3.50E+07		3.17E+05
7		279		131072		3.66E+07		3.26E+05
8		291		131072		3.81E+07		3.27E+05
9		303		131072		3.97E+07		3.53E+05
10		315		131072		4.13E+07		3.57E+05
11		327		131072		4.29E+07		3.73E+05
12		340		131072		4.46E+07		3.98E+05
13		352		131072		4.61E+07		3.96E+05
14		364		131072		4.77E+07		4.33E+05
15		367		131072		4.81E+07		4.41E+05

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		388		65536		2.54E+07		2.46E+05
1		412		65536		2.70E+07		2.47E+05
2		436		65536		2.86E+07		2.49E+05
3		461		65536		3.02E+07		2.56E+05
4		485		65536		3.18E+07		2.60E+05
5		509		65536		3.34E+07		2.95E+05
6		533		65536		3.49E+07		3.02E+05
7		558		65536		3.66E+07		3.04E+05
8		582		65536		3.81E+07		3.07E+05
9		606		65536		3.97E+07		3.11E+05
10		630		65536		4.13E+07		3.07E+05
11		654		65536		4.29E+07		3.15E+05
12		679		65536		4.45E+07		3.22E+05
13		703		65536		4.61E+07		3.22E+05
14		727		65536		4.76E+07		3.28E+05
15		733		65536		4.80E+07		3.28E+05

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	775	32768	2.54E+07	1.90E+05
1	824	32768	2.70E+07	2.18E+05
2	872	32768	2.86E+07	2.30E+05
3	921	32768	3.02E+07	2.34E+05
4	969	32768	3.18E+07	2.30E+05
5	1018	32768	3.34E+07	2.37E+05
6	1066	32768	3.49E+07	2.44E+05
7	1115	32768	3.65E+07	2.46E+05
8	1163	32768	3.81E+07	2.43E+05
9	1211	32768	3.97E+07	2.42E+05
10	1260	32768	4.13E+07	2.49E+05
11	1308	32768	4.29E+07	2.53E+05
12	1357	32768	4.45E+07	2.57E+05
13	1405	32768	4.60E+07	2.58E+05
14	1454	32768	4.76E+07	2.59E+05
15	1465	32768	4.80E+07	2.55E+05

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	1550	16384	2.54E+07	1.36E+05
1	1647	16384	2.70E+07	1.33E+05
2	1744	16384	2.86E+07	1.34E+05
3	1841	16384	3.02E+07	1.35E+05
4	1938	16384	3.18E+07	1.74E+05
5	2035	16384	3.33E+07	1.82E+05
6	2132	16384	3.49E+07	1.77E+05
7	2229	16384	3.65E+07	1.83E+05
8	2325	16384	3.81E+07	1.88E+05
9	2422	16384	3.97E+07	1.80E+05
10	2519	16384	4.13E+07	1.81E+05
11	2616	16384	4.29E+07	1.82E+05
12	2713	16384	4.44E+07	1.83E+05
13	2810	16384	4.60E+07	1.84E+05
14	2907	16384	4.76E+07	1.84E+05
15	2930	16384	4.80E+07	1.91E+05

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		3100		8192		2.54E+07		1.31E+05
1		3294		8192		2.70E+07		1.21E+05
2		3488		8192		2.86E+07		1.34E+05
3		3682		8192		3.02E+07		1.25E+05
4		3875		8192		3.17E+07		1.36E+05
5		4096		8192		3.36E+07		1.25E+05
6		4263		8192		3.49E+07		1.34E+05
7		4457		8192		3.65E+07		1.39E+05
8		4650		8192		3.81E+07		1.27E+05
9		4844		8192		3.97E+07		1.35E+05
10		5038		8192		4.13E+07		1.41E+05
11		5232		8192		4.29E+07		1.41E+05
12		5425		8192		4.44E+07		1.42E+05
13		5619		8192		4.60E+07		1.43E+05
14		5813		8192		4.76E+07		1.43E+05
15		5859		8192		4.80E+07		1.44E+05

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		6200		4096		2.54E+07		7.21E+04
1		6588		4096		2.70E+07		7.06E+04
2		6975		4096		2.86E+07		7.27E+04
3		7363		4096		3.02E+07		6.86E+04
4		7750		4096		3.17E+07		6.91E+04
5		8138		4096		3.33E+07		7.03E+04
6		8525		4096		3.49E+07		7.41E+04
7		8913		4096		3.65E+07		7.43E+04
8		9300		4096		3.81E+07		7.00E+04
9		9688		4096		3.97E+07		6.60E+04
10		10075		4096		4.13E+07		6.95E+04
11		10463		4096		4.29E+07		7.14E+04
12		10850		4096		4.44E+07		7.23E+04
13		11238		4096		4.60E+07		7.53E+04
14		11625		4096		4.76E+07		7.55E+04
15		11717		4096		4.80E+07		7.57E+04

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		12400		2048		2.54E+07		3.56E+04
1		13175		2048		2.70E+07		3.41E+04
2		13950		2048		2.86E+07		3.46E+04
3		14725		2048		3.02E+07		3.65E+04
4		15500		2048		3.17E+07		3.74E+04
5		16275		2048		3.33E+07		3.84E+04
6		17050		2048		3.49E+07		3.85E+04
7		17825		2048		3.65E+07		3.84E+04
8		18600		2048		3.81E+07		3.78E+04
9		19375		2048		3.97E+07		3.86E+04
10		20150		2048		4.13E+07		3.87E+04
11		20925		2048		4.29E+07		3.86E+04
12		21700		2048		4.44E+07		3.98E+04
13		22475		2048		4.60E+07		3.88E+04
14		23250		2048		4.76E+07		3.88E+04
15		23433		2048		4.80E+07		3.89E+04

Block Write Test 2:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache			X	

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	194	131072	2.54E+07	3.28E+05
1	206	131072	2.70E+07	3.23E+05
2	218	131072	2.86E+07	3.43E+05
3	231	131072	3.03E+07	3.41E+05
4	243	131072	3.19E+07	3.72E+05
5	255	131072	3.34E+07	3.78E+05
6	267	131072	3.50E+07	4.00E+05
7	279	131072	3.66E+07	4.00E+05
8	291	131072	3.81E+07	4.13E+05
9	303	131072	3.97E+07	4.23E+05
10	315	131072	4.13E+07	4.37E+05
11	327	131072	4.29E+07	4.52E+05
12	340	131072	4.46E+07	4.59E+05
13	352	131072	4.61E+07	4.62E+05
14	364	131072	4.77E+07	5.00E+05
15	367	131072	4.81E+07	5.11E+05

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	388	65536	2.54E+07	3.81E+05
1	412	65536	2.70E+07	3.60E+05
2	436	65536	2.86E+07	3.99E+05
3	461	65536	3.02E+07	3.92E+05
4	485	65536	3.18E+07	4.21E+05
5	509	65536	3.34E+07	4.22E+05
6	533	65536	3.49E+07	4.41E+05
7	558	65536	3.66E+07	4.61E+05
8	582	65536	3.81E+07	4.70E+05
9	606	65536	3.97E+07	4.75E+05
10	630	65536	4.13E+07	4.89E+05
11	654	65536	4.29E+07	4.98E+05
12	679	65536	4.45E+07	5.15E+05
13	703	65536	4.61E+07	5.17E+05
14	727	65536	4.76E+07	5.41E+05
15	733	65536	4.80E+07	5.60E+05

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	775	32768	2.54E+07	3.93E+05
1	824	32768	2.70E+07	4.03E+05
2	872	32768	2.86E+07	4.17E+05
3	921	32768	3.02E+07	4.33E+05
4	969	32768	3.18E+07	4.57E+05
5	1018	32768	3.34E+07	4.69E+05
6	1066	32768	3.49E+07	4.85E+05
7	1115	32768	3.65E+07	4.88E+05
8	1163	32768	3.81E+07	5.01E+05
9	1211	32768	3.97E+07	5.33E+05
10	1260	32768	4.13E+07	5.43E+05
11	1308	32768	4.29E+07	5.53E+05
12	1357	32768	4.45E+07	5.52E+05
13	1405	32768	4.60E+07	5.67E+05
14	1454	32768	4.76E+07	6.07E+05
15	1465	32768	4.80E+07	6.16E+05

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate
0	1550	16384	2.54E+07	3.54E+05
1	1647	16384	2.70E+07	3.58E+05
2	1744	16384	2.86E+07	3.78E+05
3	1841	16384	3.02E+07	3.84E+05
4	1938	16384	3.18E+07	4.00E+05
5	2035	16384	3.33E+07	4.30E+05
6	2132	16384	3.49E+07	4.33E+05
7	2229	16384	3.65E+07	4.41E+05
8	2325	16384	3.81E+07	4.44E+05
9	2422	16384	3.97E+07	4.54E+05
10	2519	16384	4.13E+07	4.70E+05
11	2616	16384	4.29E+07	4.80E+05
12	2713	16384	4.44E+07	4.77E+05
13	2810	16384	4.60E+07	4.80E+05
14	2907	16384	4.76E+07	4.99E+05
15	2930	16384	4.80E+07	5.22E+05

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		3100		8192		2.54E+07		2.94E+05
1		3294		8192		2.70E+07		3.04E+05
2		3488		8192		2.86E+07		3.15E+05
3		3682		8192		3.02E+07		3.26E+05
4		3875		8192		3.17E+07		3.35E+05
5		4096		8192		3.36E+07		3.52E+05
6		4263		8192		3.49E+07		3.64E+05
7		4457		8192		3.65E+07		3.66E+05
8		4650		8192		3.81E+07		3.73E+05
9		4844		8192		3.97E+07		3.88E+05
10		5038		8192		4.13E+07		3.98E+05
11		5232		8192		4.29E+07		4.05E+05
12		5425		8192		4.44E+07		4.02E+05
13		5619		8192		4.60E+07		4.17E+05
14		5813		8192		4.76E+07		4.26E+05
15		5859		8192		4.80E+07		4.36E+05

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		6200		4096		2.54E+07		2.27E+05
1		6588		4096		2.70E+07		2.28E+05
2		6975		4096		2.86E+07		2.36E+05
3		7363		4096		3.02E+07		2.35E+05
4		7750		4096		3.17E+07		2.57E+05
5		8138		4096		3.33E+07		2.60E+05
6		8525		4096		3.49E+07		2.63E+05
7		8913		4096		3.65E+07		2.60E+05
8		9300		4096		3.81E+07		2.61E+05
9		9688		4096		3.97E+07		2.71E+05
10		10075		4096		4.13E+07		2.70E+05
11		10463		4096		4.29E+07		2.76E+05
12		10850		4096		4.44E+07		2.73E+05
13		11238		4096		4.60E+07		2.67E+05
14		11625		4096		4.76E+07		2.91E+05
15		11717		4096		4.80E+07		2.88E+05

Block Transfer Rate

Band Number		Number of CDBs		Block Size		Total Stored		Transfer Rate
0		12400		2048		2.54E+07		1.66E+05
1		13175		2048		2.70E+07		1.66E+05
2		13950		2048		2.86E+07		1.66E+05
3		14725		2048		3.02E+07		1.67E+05
4		15500		2048		3.17E+07		1.71E+05
5		16275		2048		3.33E+07		1.72E+05
6		17050		2048		3.49E+07		1.74E+05
7		17825		2048		3.65E+07		1.74E+05
8		18600		2048		3.81E+07		1.73E+05
9		19375		2048		3.97E+07		1.75E+05
10		20150		2048		4.13E+07		1.75E+05
11		20925		2048		4.29E+07		1.78E+05
12		21700		2048		4.44E+07		1.75E+05
13		22475		2048		4.60E+07		1.75E+05
14		23250		2048		4.76E+07		1.77E+05
15		23433		2048		4.80E+07		1.77E+05

Read Cache Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode				
Read Cache			X	X

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		194	131072	2.54E+07	516.15		643.98
1		206	131072	2.70E+07	505.42		630.96
2		218	131072	2.86E+07	535.89		655.49
3		231	131072	3.03E+07	544.83		690.20
4		243	131072	3.19E+07	574.30		688.14
5		255	131072	3.34E+07	583.17		696.69
6		267	131072	3.50E+07	611.82		719.34
7		279	131072	3.66E+07	613.40		705.21
8		291	131072	3.81E+07	623.29		723.82
9		303	131072	3.97E+07	604.49		707.48
10		315	131072	4.13E+07	616.89		703.79
11		327	131072	4.29E+07	645.83		709.54
12		340	131072	4.46E+07	647.91		737.00
13		352	131072	4.61E+07	649.97		736.93
14		364	131072	4.77E+07	693.02		757.47
15		367	131072	4.81E+07	721.15		780.33

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		388	65536	2.54E+07	396.55		540.77
1		412	65536	2.70E+07	401.77		558.17
2		436	65536	2.86E+07	399.37		580.61
3		461	65536	3.02E+07	400.54		598.22
4		485	65536	3.18E+07	438.42		618.33
5		509	65536	3.34E+07	430.39		619.79
6		533	65536	3.49E+07	461.41		597.83
7		558	65536	3.66E+07	446.23		591.55
8		582	65536	3.81E+07	442.06		634.44
9		606	65536	3.97E+07	484.98		604.58
10		630	65536	4.13E+07	458.55		629.02
11		654	65536	4.29E+07	470.40		617.53
12		679	65536	4.45E+07	513.42		642.27
13		703	65536	4.61E+07	475.45		649.14
14		727	65536	4.76E+07	485.73		677.07
15		733	65536	4.80E+07	532.18		671.51

Read Cache Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		775	32768	2.54E+07	342.31		462.60
1		824	32768	2.70E+07	309.74		416.03
2		872	32768	2.86E+07	360.05		452.77
3		921	32768	3.02E+07	367.76		458.99
4		969	32768	3.18E+07	329.59		482.92
5		1018	32768	3.34E+07	382.62		502.25
6		1066	32768	3.49E+07	392.59		486.00
7		1115	32768	3.65E+07	395.61		532.94
8		1163	32768	3.81E+07	400.00		509.81
9		1211	32768	3.97E+07	380.74		467.57
10		1260	32768	4.13E+07	413.84		516.59
11		1308	32768	4.29E+07	414.37		467.51
12		1357	32768	4.45E+07	427.15		532.35
13		1405	32768	4.60E+07	431.02		498.23
14		1454	32768	4.76E+07	430.26		563.98
15		1465	32768	4.80E+07	436.13		558.96

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		1550	16384	2.54E+07	286.14		295.70
1		1647	16384	2.70E+07	295.79		311.12
2		1744	16384	2.86E+07	282.89		292.31
3		1841	16384	3.02E+07	305.59		332.50
4		1938	16384	3.18E+07	293.41		315.57
5		2035	16384	3.33E+07	301.68		335.84
6		2132	16384	3.49E+07	301.05		345.02
7		2229	16384	3.65E+07	329.43		339.08
8		2325	16384	3.81E+07	334.62		340.53
9		2422	16384	3.97E+07	330.45		382.40
10		2519	16384	4.13E+07	312.77		347.60
11		2616	16384	4.29E+07	344.04		391.80
12		2713	16384	4.44E+07	375.08		413.13
13		2810	16384	4.60E+07	324.06		366.75
14		2907	16384	4.76E+07	325.83		362.02
15		2930	16384	4.80E+07	331.07		340.45

Read Cache Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		3100	8192	2.54E+07				221.77
1		3294	8192	2.70E+07				186.68
2		3488	8192	2.86E+07				216.83
3		3682	8192	3.02E+07				203.59
4		3875	8192	3.17E+07				223.34
5		4096	8192	3.36E+07				239.35
6		4263	8192	3.49E+07				241.89
7		4457	8192	3.65E+07				244.40
8		4650	8192	3.81E+07				247.92
9		4844	8192	3.97E+07				249.48
10		5038	8192	4.13E+07				251.73
11		5232	8192	4.29E+07				251.18
12		5425	8192	4.44E+07				255.88
13		5619	8192	4.60E+07				257.13
14		5813	8192	4.76E+07				258.36
15		5859	8192	4.80E+07				260.66

Block Transfer Rate


Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		6200	4096	2.54E+07				131.22
1		6588	4096	2.70E+07				132.65
2		6975	4096	2.86E+07				133.74
3		7363	4096	3.02E+07				135.03
4		7750	4096	3.17E+07				135.84
5		8138	4096	3.33E+07				136.78
6		8525	4096	3.49E+07				137.87
7		8913	4096	3.65E+07				138.31
8		9300	4096	3.81E+07				139.04
9		9688	4096	3.97E+07				139.71
10		10075	4096	4.13E+07				141.89
11		10463	4096	4.29E+07				141.45
12		10850	4096	4.44E+07				142.01
13		11238	4096	4.60E+07				142.43
14		11625	4096	4.76E+07				142.38
15		11717	4096	4.80E+07				142.67

Read Cache Test 1 (Continued)


Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		12400	2048	2.54E+07			76.97
1		13175	2048	2.70E+07			72.37
2		13950	2048	2.86E+07			72.89
3		14725	2048	3.02E+07			73.26
4		15500	2048	3.17E+07			73.51
5		16275	2048	3.33E+07			73.86
6		17050	2048	3.49E+07			74.13
7		17825	2048	3.65E+07			74.32
8		18600	2048	3.81E+07			74.54
9		19375	2048	3.97E+07			74.72
10		20150	2048	4.13E+07			74.94
11		20925	2048	4.29E+07			74.86
12		21700	2048	4.44E+07			75.30
13		22475	2048	4.60E+07			75.43
14		23250	2048	4.76E+07			75.53
15		23433	2048	4.80E+07			75.74


Hewlett-Packard C1716T Average Read Rate




Hewlett-Packard C1716T Average Write Data Rates




Hewlett-Packard C1716T Write Cache Testing




Hewlett-Packard C1716T Band Testing (Cache Disabled)




Hewlett-Packard CP C1716T Band Testing (Cache Enabled)




Hewlett-Packard C1716T 128k Write




Hewlett-Packard C1716T 64k Write




Hewlett Packard C1716T 32k Write





Hewlett-Packard C1716T 16k Write



Hewlett-Packard C1716T 8k Write





Optical Drive Description:

Hitachi OD152S-1 Standard, 5.25 inch, half height

Interface: Single Ended Cache: 512 kB Rotational Speed: 3000

Power Requirements:

	+5 VDC	+12 VDC
Tolerance	+/- 5%	+/- 5%
Ripple	100 mV PP (max)	100 mV PP (max)
Current:	Typical	0.6 A
	Maximum	2.3 A
1.5 A		
1.7 A		

2 Pass Write: N/A

3 Pass Write: Erase/Write/Verify

Media Description:

Maxell OC-152-S1 Format: ZCAV Type: 512 Byte/sector Formatted Capacity: 1.738 GB

Read Test:

Read Cache Enabled

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
6634	128	849152	1614.63
13268	64	849152	1300.80
26535	32	849120	919.11
53069	16	849104	630.49
106137	8	849096	529.54
212273	4	849092	290.33
424545	2	849090	178.55

Read Cache Disabled

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
6634	128	849152	1428.18
13268	64	849152	1122.98
26535	32	849120	814.17
53069	16	849104	524.50
106137	8	849096	331.29
212273	4	849092	185.41
424545	2	849090	95.02

Write Test:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode	X			
Write Cache				X
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
6634	128	849152	483.29
13268	64	849152	393.31
26535	32	849120	305.69
53069	16	849104	211.98
106137	8	849096	119.86
212273	4	849092	62.50
424545	2	849090	31.71

Write Cache Test:

Conditions	3 Pass	2 Pass	Read	Write
Write Mode	X			
Test Mode				X
Burst Size: unlimited				X

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Cache Disabled	Cache Enabled
2	131072	262144		360.56
4	131072	524288		358.04
6	131072	786432		349.09
4	65536	262144		312.20
8	65536	524288		299.42
12	65536	786432		303.56
8	32768	262144		220.69
16	32768	524288		222.61
24	32768	786432		218.80
16	16384	262144		166.23
32	16384	524288		166.78
48	16384	786432		170.67
32	8192	262144		111.30
64	8192	524288		112.28
96	8192	786432		111.79
64	4096	262144		58.31
128	4096	524288		58.25
192	4096	786432		58.54
128	2048	262144		31.30
256	2048	524288		31.72
384	2048	786432		31.55

Block Read Test:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode				
Read Cache			X	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		79	131072	1.04E+07	1173.09		1057.74
1		80	131072	1.05E+07	1122.81		1053.50
2		82	131072	1.07E+07	1217.63		1097.91
3		83	131072	1.09E+07	1208.65		1104.37
4		85	131072	1.11E+07	1269.55		1089.09
5		86	131072	1.13E+07	1228.57		1083.46
6		88	131072	1.15E+07	1250.17		1164.84
7		90	131072	1.18E+07	1271.52		1185.19
8		91	131072	1.19E+07	1292.79		1164.80
9		93	131072	1.22E+07	1298.15		1204.86
10		94	131072	1.23E+07	1273.23		1216.58
11		96	131072	1.26E+07	1315.63		1215.43
12		97	131072	1.27E+07	1256.68		1203.10
13		99	131072	1.30E+07	1381.90		1333.89
14		101	131072	1.32E+07	1425.36		1266.21
15		102	131072	1.34E+07	1406.90		1291.39
16		104	131072	1.36E+07	1369.55		1261.80
17		105	131072	1.38E+07	1429.79		1254.90
18		107	131072	1.40E+07	1441.68		1265.80
19		109	131072	1.43E+07	1459.41		1344.12
20		110	131072	1.44E+07	1525.46		1314.66
21		112	131072	1.47E+07	1536.55		1324.95
22		113	131072	1.48E+07	1605.33		1408.37
23		115	131072	1.51E+07	1739.95		1346.75
24		116	131072	1.52E+07	1649.78		1358.46
25		118	131072	1.55E+07	1706.67		1493.97
26		120	131072	1.57E+07	1634.04		1413.06
27		121	131072	1.59E+07	1638.94		1484.95
28		123	131072	1.61E+07	1601.63		1470.03
29		124	131072	1.63E+07	1606.48		1402.12
30		126	131072	1.65E+07	1697.68		1537.46
31		127	131072	1.66E+07	1731.20		1542.32
32		129	131072	1.69E+07	1747.30		1526.06
33		131	131072	1.72E+07	1816.69		1453.03
34		132	131072	1.73E+07	1653.23		1451.55
35		134	131072	1.76E+07	1652.41		1553.62
36		135	131072	1.77E+07	1874.19		1565.22
37		137	131072	1.80E+07	1764.19		1572.74
38		138	131072	1.81E+07	1891.22		1591.35
39		140	131072	1.84E+07	1802.82		1560.98
40		142	131072	1.86E+07	1788.98		1546.89

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
				Enabled		Disabled
41	143	131072	1.87E+07	1892.87		1556.46
42	145	131072	1.90E+07	1760.91		1732.96
43	146	131072	1.91E+07	1910.84		1676.05
44	148	131072	1.94E+07	1777.11		1604.06
45	149	131072	1.95E+07	2119.11		1638.49
46	151	131072	1.98E+07	2119.30		1621.48
47	153	131072	2.01E+07	1958.40		1756.41
48	154	131072	2.02E+07	2005.29		1669.09
49	156	131072	2.04E+07	1873.17		1722.86
50	157	131072	2.06E+07	1733.91		1701.61
51	159	131072	2.08E+07	1825.29		1765.13
52	160	131072	2.10E+07	1846.71		1627.98
53	162	131072	2.12E+07	1907.64		1851.43
54	164	131072	2.15E+07	1856.06		1729.16

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
				Enabled		Disabled
0	79	65536	5.18E+06	1135.37		914.29
1	80	65536	5.24E+06	1187.94		914.29
2	82	65536	5.37E+06	1172.14		884.07
3	83	65536	5.44E+06	962.32		911.93
4	85	65536	5.57E+06	1215.28		947.11
5	86	65536	5.64E+06	1145.47		968.16
6	88	65536	5.77E+06	1073.79		911.33
7	90	65536	5.90E+06	1218.72		957.06
8	91	65536	5.96E+06	1109.33		1005.00
9	93	65536	6.09E+06	1105.51		1006.80
10	94	65536	6.16E+06	1423.91		908.76
11	96	65536	6.29E+06	1153.21		993.01
12	97	65536	6.36E+06	1124.54		1046.88
13	99	65536	6.49E+06	1281.29		994.66
14	101	65536	6.62E+06	1239.31		1036.58
15	102	65536	6.68E+06	1170.94		998.17
16	104	65536	6.82E+06	1236.98		1086.79
17	105	65536	6.88E+06	1236.43		1002.24
18	107	65536	7.01E+06	1216.06		1084.49
19	109	65536	7.14E+06	1085.00		1062.59
20	110	65536	7.21E+06	1187.18		1081.41
21	112	65536	7.34E+06	1175.62		1069.87
22	113	65536	7.41E+06	1265.44		1101.60
23	115	65536	7.54E+06	1169.67		1084.83
24	116	65536	7.60E+06	1072.83		1081.43
25	118	65536	7.73E+06	1280.00		1065.16
26	120	65536	7.86E+06	1358.44		1156.16
27	121	65536	7.93E+06	1508.08		1105.50

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
28		123	65536	8.06E+06		1328.81		1098.04
29		124	65536	8.13E+06		1307.41		1111.48
30		126	65536	8.26E+06		1265.88		1120.22
31		127	65536	8.32E+06		1253.35		1213.13
32		129	65536	8.45E+06		1343.79		1143.02
33		131	65536	8.59E+06		1363.59		1201.73
34		132	65536	8.65E+06		1235.09		1192.38
35		134	65536	8.78E+06		1358.35		1047.70
36		135	65536	8.85E+06		1328.21		1169.15
37		137	65536	8.98E+06		1353.37		1233.03
38		138	65536	9.04E+06		1380.00		1169.80
39		140	65536	9.18E+06		1426.20		1181.73
40		142	65536	9.31E+06		1335.69		1303.96
41		143	65536	9.37E+06		1460.81		1221.08
42		145	65536	9.50E+06		1451.81		1228.97
43		146	65536	9.57E+06		1479.65		1328.22
44		148	65536	9.70E+06		1462.43		1278.27
45		149	65536	9.76E+06		1400.29		1214.00
46		151	65536	9.90E+06		1407.72		1270.74
47		153	65536	1.00E+07		1439.53		1356.50
48		154	65536	1.01E+07		1401.99		1295.14
49		156	65536	1.02E+07		1357.71		1250.25
50		157	65536	1.03E+07		1499.70		1231.37
51		159	65536	1.04E+07		1333.86		1353.44
52		160	65536	1.05E+07		1376.34		1331.60
53		162	65536	1.06E+07		1367.81		1154.57
54		164	65536	1.07E+07		1360.73		1270.67

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		79	32768	2.59E+06		799.36		706.85
1		80	32768	2.62E+06		826.56		717.36
2		82	32768	2.69E+06		785.50		717.24
3		83	32768	2.72E+06		805.46		721.74
4		85	32768	2.79E+06		933.22		639.24
5		86	32768	2.82E+06		824.57		691.02
6		88	32768	2.88E+06		821.65		684.04
7		90	32768	2.95E+06		812.52		673.19
8		91	32768	2.98E+06		859.81		755.27
9		93	32768	3.05E+06		941.18		713.68
10		94	32768	3.08E+06		859.43		706.52
11		96	32768	3.15E+06		859.03		786.62
12		97	32768	3.18E+06		893.24		787.32
13		99	32768	3.24E+06		946.82		804.58
14		101	32768	3.31E+06		993.96		803.00
15		102	32768	3.34E+06		948.58		785.05
16		104	32768	3.41E+06		913.66		720.00
17		105	32768	3.44E+06		883.63		810.13
18		107	32768	3.51E+06		849.88		816.29
19		109	32768	3.57E+06		808.40		729.26
20		110	32768	3.60E+06		916.97		716.37
21		112	32768	3.67E+06		909.32		839.13
22		113	32768	3.70E+06		890.86		839.56
23		115	32768	3.77E+06		988.27		775.86
24		116	32768	3.80E+06		951.79		757.16
25		118	32768	3.87E+06		869.36		752.00
26		120	32768	3.93E+06		893.68		834.55
27		121	32768	3.96E+06		910.78		832.31
28		123	32768	4.03E+06		924.82		713.87
29		124	32768	4.06E+06		905.94		758.34
30		126	32768	4.13E+06		753.82		859.96
31		127	32768	4.16E+06		948.98		878.23
32		129	32768	4.23E+06		953.51		902.25
33		131	32768	4.29E+06		998.32		906.09
34		132	32768	4.33E+06		905.69		869.73
35		134	32768	4.39E+06		989.33		916.00
36		135	32768	4.42E+06		977.38		928.03
37		137	32768	4.49E+06		976.09		924.44
38		138	32768	4.52E+06		1008.22		929.68
39		140	32768	4.59E+06		868.91		810.53
40		142	32768	4.65E+06		994.50		797.18
41		143	32768	4.69E+06		984.48		792.02
42		145	32768	4.75E+06		1000.22		880.08
43		146	32768	4.78E+06		872.46		929.75
44		148	32768	4.85E+06		984.87		919.18
45		149	32768	4.88E+06		998.01		863.77
46		151	32768	4.95E+06		919.73		802.33

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled	Disabled	
47		153	32768	5.01E+06		1013.94		834.96
48		154	32768	5.05E+06		978.62		810.88
49		156	32768	5.11E+06		911.56		876.19
50		157	32768	5.14E+06		965.23		970.35
51		159	32768	5.21E+06		932.78		959.24
52		160	32768	5.24E+06		1002.45		939.45
53		162	32768	5.31E+06		1021.92		885.16
54		164	32768	5.37E+06		1003.65		864.54

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled	Disabled	
0		79	16384	1.29E+06		515.73		479.16
1		80	16384	1.31E+06		601.18		465.69
2		82	16384	1.34E+06		568.62		464.08
3		83	16384	1.36E+06		505.38		431.04
4		85	16384	1.39E+06		521.00		421.68
5		86	16384	1.41E+06		544.68		503.57
6		88	16384	1.44E+06		563.90		475.05
7		90	16384	1.47E+06		547.93		509.06
8		91	16384	1.49E+06		682.09		516.96
9		93	16384	1.52E+06		681.11		519.96
10		94	16384	1.54E+06		708.07		524.72
11		96	16384	1.57E+06		647.11		527.35
12		97	16384	1.59E+06		602.72		529.24
13		99	16384	1.62E+06		606.34		534.69
14		101	16384	1.65E+06		613.09		530.91
15		102	16384	1.67E+06		600.18		537.52
16		104	16384	1.70E+06		627.41		540.30
17		105	16384	1.72E+06		596.09		540.64
18		107	16384	1.75E+06		639.70		546.59
19		109	16384	1.79E+06		636.32		551.39
20		110	16384	1.80E+06		591.32		549.41
21		112	16384	1.84E+06		702.86		552.80
22		113	16384	1.85E+06		693.51		434.31
23		115	16384	1.88E+06		664.85		478.59
24		116	16384	1.90E+06		707.72		512.88
25		118	16384	1.93E+06		705.44		559.94
26		120	16384	1.97E+06		599.61		550.87
27		121	16384	1.98E+06		617.60		555.60
28		123	16384	2.02E+06		656.38		536.63
29		124	16384	2.03E+06		654.67		486.83
30		126	16384	2.06E+06		625.65		486.16
31		127	16384	2.08E+06		665.14		485.25
32		129	16384	2.11E+06		688.49		443.58
33		131	16384	2.15E+06		654.97		494.68

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
34		132	16384	2.16E+06		633.86		507.01
35		134	16384	2.20E+06		655.24		579.35
36		135	16384	2.21E+06		623.70		583.24
37		137	16384	2.24E+06		623.21		580.90
38		138	16384	2.26E+06		709.97		582.59
39		140	16384	2.29E+06		716.25		564.53
40		142	16384	2.33E+06		707.28		517.15
41		143	16384	2.34E+06		628.98		512.54
42		145	16384	2.38E+06		652.67		499.43
43		146	16384	2.39E+06		655.16		505.88
44		148	16384	2.42E+06		691.24		471.72
45		149	16384	2.44E+06		659.02		469.87
46		151	16384	2.47E+06		653.56		507.23
47		153	16384	2.51E+06		628.53		515.95
48		154	16384	2.52E+06		645.46		591.35
49		156	16384	2.56E+06		660.20		601.09
50		157	16384	2.57E+06		654.07		596.38
51		159	16384	2.61E+06		660.63		606.52
52		160	16384	2.62E+06		694.24		608.26
53		162	16384	2.65E+06		673.88		604.56
54		164	16384	2.69E+06		630.43		592.17

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		79	8192	6.47E+05		445.24		303.95
1		80	8192	6.55E+05		429.84		305.94
2		82	8192	6.72E+05		457.34		306.88
3		83	8192	6.80E+05		452.37		308.46
4		85	8192	6.96E+05		461.88		309.42
5		86	8192	7.05E+05		491.21		311.22
6		88	8192	7.21E+05		459.72		312.55
7		90	8192	7.37E+05		487.52		313.75
8		91	8192	7.45E+05		487.25		314.20
9		93	8192	7.62E+05		459.59		315.43
10		94	8192	7.70E+05		490.45		315.22
11		96	8192	7.86E+05		502.10		317.71
12		97	8192	7.95E+05		488.24		318.85
13		99	8192	8.11E+05		485.91		320.49
14		101	8192	8.27E+05		482.47		321.16
15		102	8192	8.36E+05		495.02		321.74
16		104	8192	8.52E+05		515.56		322.30
17		105	8192	8.60E+05		505.99		322.46
18		107	8192	8.77E+05		492.20		324.23
19		109	8192	8.93E+05		517.08		326.23
20		110	8192	9.01E+05		501.86		325.87

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
21		112	8192	9.18E+05		579.71		327.17
22		113	8192	9.26E+05		541.74		328.00
23		115	8192	9.42E+05		563.20		328.88
24		116	8192	9.50E+05		544.79		329.48
25		118	8192	9.67E+05		599.20		329.54
26		120	8192	9.83E+05		445.35		329.51
27		121	8192	9.91E+05		556.60		331.54
28		123	8192	1.01E+06		517.71		331.85
29		124	8192	1.02E+06		541.26		333.15
30		126	8192	1.03E+06		550.34		334.33
31		127	8192	1.04E+06		513.53		333.16
32		129	8192	1.06E+06		506.72		334.58
33		131	8192	1.07E+06		554.01		335.65
34		132	8192	1.08E+06		459.95		334.55
35		134	8192	1.10E+06		543.01		336.74
36		135	8192	1.11E+06		540.77		338.22
37		137	8192	1.12E+06		526.10		338.10
38		138	8192	1.13E+06		519.29		337.66
39		140	8192	1.15E+06		521.80		339.73
40		142	8192	1.16E+06		511.58		340.12
41		143	8192	1.17E+06		531.08		340.30
42		145	8192	1.19E+06		565.86		341.67
43		146	8192	1.20E+06		572.16		341.21
44		148	8192	1.21E+06		557.45		341.74
45		149	8192	1.22E+06		597.68		342.41
46		151	8192	1.24E+06		571.53		342.98
47		153	8192	1.25E+06		582.99		343.54
48		154	8192	1.26E+06		565.78		344.70
49		156	8192	1.28E+06		544.89		343.57
50		157	8192	1.29E+06		558.80		345.06
51		159	8192	1.30E+06		572.40		345.43
52		160	8192	1.31E+06		586.15		346.24
53		162	8192	1.33E+06		576.69		347.04
54		164	8192	1.34E+06		598.80		346.56

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		79	4096	3.24E+05		255.69		172.74
1		80	4096	3.28E+05		260.80		173.23
2		82	4096	3.36E+05		260.46		173.85
3		83	4096	3.40E+05		265.96		174.09
4		85	4096	3.48E+05		270.57		174.54
5		86	4096	3.52E+05		280.32		174.87
6		88	4096	3.60E+05		275.72		174.85
7		90	4096	3.69E+05		272.92		178.14
8		91	4096	3.73E+05		280.08		194.51
9		93	4096	3.81E+05		273.24		221.03
10		94	4096	3.85E+05		281.05		248.23
11		96	4096	3.93E+05		281.86		256.80
12		97	4096	3.97E+05		273.30		261.94
13		99	4096	4.06E+05		270.93		227.96
14		101	4096	4.14E+05		273.00		200.34
15		102	4096	4.18E+05		266.37		179.04
16		104	4096	4.26E+05		274.19		179.14
17		105	4096	4.30E+05		296.27		177.75
18		107	4096	4.38E+05		273.79		178.18
19		109	4096	4.46E+05		277.28		178.79
20		110	4096	4.51E+05		272.55		178.84
21		112	4096	4.59E+05		279.50		179.92
22		113	4096	4.63E+05		281.05		180.11
23		115	4096	4.71E+05		280.68		180.35
24		116	4096	4.75E+05		288.30		180.55
25		118	4096	4.83E+05		289.10		180.94
26		120	4096	4.92E+05		278.49		180.84
27		121	4096	4.96E+05		288.92		181.37
28		123	4096	5.04E+05		291.46		180.89
29		124	4096	5.08E+05		285.02		181.78
30		126	4096	5.16E+05		290.97		182.24
31		127	4096	5.20E+05		289.35		181.66
32		129	4096	5.28E+05		292.72		182.27
33		131	4096	5.37E+05		297.76		182.58
34		132	4096	5.41E+05		290.73		182.16
35		134	4096	5.49E+05		298.22		182.94
36		135	4096	5.53E+05		297.70		183.26
37		137	4096	5.61E+05		294.02		183.07
38		138	4096	5.65E+05		300.43		183.06
39		140	4096	5.73E+05		298.55		183.57
40		142	4096	5.82E+05		298.91		183.80
41		143	4096	5.86E+05		300.87		183.98
42		145	4096	5.94E+05		298.22		184.18
43		146	4096	5.98E+05		301.76		184.06
44		148	4096	6.06E+05		306.93		184.06
45		149	4096	6.10E+05		303.97		184.27
46		151	4096	6.18E+05		302.96		184.63

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
47		153	4096	6.27E+05		315.63		184.41
48		154	4096	6.31E+05		305.61		185.01
49		156	4096	6.39E+05		308.43		184.98
50		157	4096	6.43E+05		308.14		185.11
51		159	4096	6.51E+05		309.42		185.24
52		160	4096	6.55E+05		308.10		185.10
53		162	4096	6.64E+05		311.96		185.65
54		164	4096	6.72E+05		308.83		185.62

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		79	2048	1.62E+05		157.47		92.58
1		80	2048	1.64E+05		158.31		92.80
2		82	2048	1.68E+05		165.24		93.00
3		83	2048	1.70E+05		161.79		93.02
4		85	2048	1.74E+05		163.38		93.18
5		86	2048	1.76E+05		165.48		93.29
6		88	2048	1.80E+05		168.18		93.24
7		90	2048	1.84E+05		166.42		93.56
8		91	2048	1.86E+05		170.58		93.45
9		93	2048	1.90E+05		170.97		93.66
10		94	2048	1.93E+05		175.67		93.79
11		96	2048	1.97E+05		179.20		93.79
12		97	2048	1.99E+05		173.87		93.95
13		99	2048	2.03E+05		175.64		94.00
14		101	2048	2.07E+05		173.90		94.18
15		102	2048	2.09E+05		181.79		94.14
16		104	2048	2.13E+05		176.74		94.29
17		105	2048	2.15E+05		170.31		94.18
18		107	2048	2.19E+05		180.79		94.49
19		109	2048	2.23E+05		184.15		94.60
20		110	2048	2.25E+05		183.20		94.59
21		112	2048	2.29E+05		174.84		94.70
22		113	2048	2.31E+05		180.67		94.64
23		115	2048	2.36E+05		179.75		94.71
24		116	2048	2.38E+05		173.46		94.90
25		118	2048	2.42E+05		183.37		95.02
26		120	2048	2.46E+05		179.42		94.93
27		121	2048	2.48E+05		177.10		94.99
28		123	2048	2.52E+05		185.02		95.10
29		124	2048	2.54E+05		179.03		95.22
30		126	2048	2.58E+05		177.99		95.29
31		127	2048	2.60E+05		177.63		95.19
32		129	2048	2.64E+05		176.44		95.36
33		131	2048	2.68E+05		176.84		95.44

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
34		132	2048	2.70E+05		169.65		95.20
35		134	2048	2.74E+05		178.20		95.48
36		135	2048	2.76E+05		178.78		95.50
37		137	2048	2.81E+05		180.05		95.60
38		138	2048	2.83E+05		180.47		95.55
39		140	2048	2.87E+05		180.11		95.70
40		142	2048	2.91E+05		181.23		95.66
41		143	2048	2.93E+05		180.86		95.72
42		145	2048	2.97E+05		181.78		95.89
43		146	2048	2.99E+05		181.59		95.85
44		148	2048	3.03E+05		181.72		95.83
45		149	2048	3.05E+05		183.08		95.94
46		151	2048	3.09E+05		182.69		95.99
47		153	2048	3.13E+05		182.15		95.95
48		154	2048	3.15E+05		184.04		96.09
49		156	2048	3.19E+05		183.10		96.12
50		157	2048	3.22E+05		185.87		96.07
51		159	2048	3.26E+05		185.97		96.18
52		160	2048	3.28E+05		184.63		96.13
53		162	2048	3.32E+05		188.43		96.21
54		164	2048	3.36E+05		187.43		96.26

Block Write Test:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode	X			
Write Cache			N/A	N/A

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		79	131072	1.04E+07		354.81	
1		80	131072	1.05E+07		352.37	
2		82	131072	1.07E+07		375.39	
3		83	131072	1.09E+07		360.87	
4		85	131072	1.11E+07		378.70	
5		86	131072	1.13E+07		380.24	
6		88	131072	1.15E+07		388.41	
7		90	131072	1.18E+07		401.81	
8		91	131072	1.19E+07		391.27	
9		93	131072	1.22E+07		393.26	
10		94	131072	1.23E+07		389.13	
11		96	131072	1.26E+07		416.68	
12		97	131072	1.27E+07		421.74	
13		99	131072	1.30E+07		443.70	
14		101	131072	1.32E+07		420.29	
15		102	131072	1.34E+07		423.76	
16		104	131072	1.36E+07		425.17	
17		105	131072	1.38E+07		456.52	
18		107	131072	1.40E+07		457.60	
19		109	131072	1.43E+07		475.69	
20		110	131072	1.44E+07		461.03	
21		112	131072	1.47E+07		464.40	
22		113	131072	1.48E+07		471.91	
23		115	131072	1.51E+07		464.50	
24		116	131072	1.52E+07		470.17	
25		118	131072	1.55E+07		486.75	
26		120	131072	1.57E+07		459.19	
27		121	131072	1.59E+07		480.40	
28		123	131072	1.61E+07		463.06	
29		124	131072	1.63E+07		473.65	
30		126	131072	1.65E+07		489.47	
31		127	131072	1.66E+07		483.67	
32		129	131072	1.69E+07		481.82	
33		131	131072	1.72E+07		479.91	
34		132	131072	1.73E+07		474.74	
35		134	131072	1.76E+07		491.74	
36		135	131072	1.77E+07		500.14	
37		137	131072	1.80E+07		488.20	
38		138	131072	1.81E+07		494.79	
39		140	131072	1.84E+07		490.69	
40		142	131072	1.86E+07		531.15	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled	Disabled	
41		143	131072	1.87E+07			533.96	
42		145	131072	1.90E+07			578.73	
43		146	131072	1.91E+07			588.60	
44		148	131072	1.94E+07			578.80	
45		149	131072	1.95E+07			572.90	
46		151	131072	1.98E+07			584.63	
47		153	131072	2.01E+07			592.20	
48		154	131072	2.02E+07			617.74	
49		156	131072	2.04E+07			590.07	
50		157	131072	2.06E+07			599.88	
51		159	131072	2.08E+07			605.35	
52		160	131072	2.10E+07			604.49	
53		162	131072	2.12E+07			601.04	
54		164	131072	2.15E+07			615.42	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled	Disabled	
0		79	65536	5.18E+06			301.83	
1		80	65536	5.24E+06			300.73	
2		82	65536	5.37E+06			311.40	
3		83	65536	5.44E+06			306.52	
4		85	65536	5.57E+06			309.12	
5		86	65536	5.64E+06			306.03	
6		88	65536	5.77E+06			310.30	
7		90	65536	5.90E+06			314.12	
8		91	65536	5.96E+06			312.78	
9		93	65536	6.09E+06			341.11	
10		94	65536	6.16E+06			350.48	
11		96	65536	6.29E+06			344.53	
12		97	65536	6.36E+06			352.63	
13		99	65536	6.49E+06			362.78	
14		101	65536	6.62E+06			347.02	
15		102	65536	6.68E+06			350.59	
16		104	65536	6.82E+06			351.03	
17		105	65536	6.88E+06			390.93	
18		107	65536	7.01E+06			393.99	
19		109	65536	7.14E+06			415.19	
20		110	65536	7.21E+06			398.64	
21		112	65536	7.34E+06			409.88	
22		113	65536	7.41E+06			411.38	
23		115	65536	7.54E+06			414.36	
24		116	65536	7.60E+06			408.36	
25		118	65536	7.73E+06			416.74	
26		120	65536	7.86E+06			401.26	
27		121	65536	7.93E+06			406.83	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)	
						Enabled	Disabled
28		123	65536	8.06E+06		413.18	
29		124	65536	8.13E+06		405.83	
30		126	65536	8.26E+06		420.85	
31		127	65536	8.32E+06		418.11	
32		129	65536	8.45E+06		411.30	
33		131	65536	8.59E+06		424.82	
34		132	65536	8.65E+06		410.10	
35		134	65536	8.78E+06		426.77	
36		135	65536	8.85E+06		433.95	
37		137	65536	8.98E+06		421.42	
38		138	65536	9.04E+06		415.43	
39		140	65536	9.18E+06		429.44	
40		142	65536	9.31E+06		425.56	
41		143	65536	9.37E+06		419.14	
42		145	65536	9.50E+06		436.74	
43		146	65536	9.57E+06		429.02	
44		148	65536	9.70E+06		418.16	
45		149	65536	9.76E+06		421.95	
46		151	65536	9.90E+06		433.36	
47		153	65536	1.00E+07		432.34	
48		154	65536	1.01E+07		429.83	
49		156	65536	1.02E+07		435.06	
50		157	65536	1.03E+07		439.74	
51		159	65536	1.04E+07		441.33	
52		160	65536	1.05E+07		441.86	
53		162	65536	1.06E+07		442.60	
54		164	65536	1.07E+07		445.56	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		79	32768	2.59E+06		220.47		
1		80	32768	2.62E+06		221.00		
2		82	32768	2.69E+06		218.67		
3		83	32768	2.72E+06		221.06		
4		85	32768	2.79E+06		253.12		
5		86	32768	2.82E+06		249.28		
6		88	32768	2.88E+06		260.54		
7		90	32768	2.95E+06		258.05		
8		91	32768	2.98E+06		252.96		
9		93	32768	3.05E+06		262.88		
10		94	32768	3.08E+06		253.25		
11		96	32768	3.15E+06		263.39		
12		97	32768	3.18E+06		265.64		
13		99	32768	3.24E+06		269.11		
14		101	32768	3.31E+06		266.11		
15		102	32768	3.34E+06		267.32		
16		104	32768	3.41E+06		264.22		
17		105	32768	3.44E+06		310.18		
18		107	32768	3.51E+06		309.82		
19		109	32768	3.57E+06		308.05		
20		110	32768	3.60E+06		325.03		
21		112	32768	3.67E+06		320.43		
22		113	32768	3.70E+06		313.13		
23		115	32768	3.77E+06		312.49		
24		116	32768	3.80E+06		323.34		
25		118	32768	3.87E+06		331.94		
26		120	32768	3.93E+06		320.87		
27		121	32768	3.96E+06		313.38		
28		123	32768	4.03E+06		305.33		
29		124	32768	4.06E+06		332.19		
30		126	32768	4.13E+06		332.73		
31		127	32768	4.16E+06		329.60		
32		129	32768	4.23E+06		333.83		
33		131	32768	4.29E+06		334.64		
34		132	32768	4.33E+06		328.73		
35		134	32768	4.39E+06		336.48		
36		135	32768	4.42E+06		340.09		
37		137	32768	4.49E+06		328.92		
38		138	32768	4.52E+06		324.89		
39		140	32768	4.59E+06		317.78		
40		142	32768	4.65E+06		327.54		
41		143	32768	4.69E+06		339.82		
42		145	32768	4.75E+06		343.32		
43		146	32768	4.78E+06		335.21		
44		148	32768	4.85E+06		327.04		
45		149	32768	4.88E+06		325.74		
46		151	32768	4.95E+06		323.22		

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
47		153	32768	5.01E+06			328.25	
48		154	32768	5.05E+06			345.81	
49		156	32768	5.11E+06			343.93	
50		157	32768	5.14E+06			340.32	
51		159	32768	5.21E+06			337.63	
52		160	32768	5.24E+06			332.36	
53		162	32768	5.31E+06			329.79	
54		164	32768	5.37E+06			326.86	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		79	16384	1.29E+06			172.56	
1		80	16384	1.31E+06			167.89	
2		82	16384	1.34E+06			166.53	
3		83	16384	1.36E+06			165.05	
4		85	16384	1.39E+06			173.81	
5		86	16384	1.41E+06			172.16	
6		88	16384	1.44E+06			167.25	
7		90	16384	1.47E+06			166.56	
8		91	16384	1.49E+06			167.04	
9		93	16384	1.52E+06			175.74	
10		94	16384	1.54E+06			176.58	
11		96	16384	1.57E+06			176.90	
12		97	16384	1.59E+06			177.17	
13		99	16384	1.62E+06			178.31	
14		101	16384	1.65E+06			177.65	
15		102	16384	1.67E+06			178.17	
16		104	16384	1.70E+06			178.45	
17		105	16384	1.72E+06			211.43	
18		107	16384	1.75E+06			230.23	
19		109	16384	1.79E+06			232.02	
20		110	16384	1.80E+06			226.65	
21		112	16384	1.84E+06			228.24	
22		113	16384	1.85E+06			215.21	
23		115	16384	1.88E+06			218.64	
24		116	16384	1.90E+06			215.41	
25		118	16384	1.93E+06			217.34	
26		120	16384	1.97E+06			230.96	
27		121	16384	1.98E+06			233.44	
28		123	16384	2.02E+06			232.58	
29		124	16384	2.03E+06			230.94	
30		126	16384	2.06E+06			231.71	
31		127	16384	2.08E+06			223.88	
32		129	16384	2.11E+06			217.16	
33		131	16384	2.15E+06			221.17	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
34		132	16384	2.16E+06			216.33	
35		134	16384	2.20E+06			218.67	
36		135	16384	2.21E+06			235.08	
37		137	16384	2.24E+06			237.34	
38		138	16384	2.26E+06			237.87	
39		140	16384	2.29E+06			237.98	
40		142	16384	2.33E+06			236.78	
41		143	16384	2.34E+06			234.77	
42		145	16384	2.38E+06			235.57	
43		146	16384	2.39E+06			229.53	
44		148	16384	2.42E+06			220.58	
45		149	16384	2.44E+06			221.74	
46		151	16384	2.47E+06			223.87	
47		153	16384	2.51E+06			222.26	
48		154	16384	2.52E+06			221.32	
49		156	16384	2.56E+06			223.73	
50		157	16384	2.57E+06			224.16	
51		159	16384	2.61E+06			237.68	
52		160	16384	2.62E+06			241.48	
53		162	16384	2.65E+06			242.39	
54		164	16384	2.69E+06			242.14	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		79	8192	6.47E+05			112.03	
1		80	8192	6.55E+05			120.11	
2		82	8192	6.72E+05			121.23	
3		83	8192	6.80E+05			120.93	
4		85	8192	6.96E+05			118.46	
5		86	8192	7.05E+05			119.80	
6		88	8192	7.21E+05			120.89	
7		90	8192	7.37E+05			122.10	
8		91	8192	7.45E+05			122.11	
9		93	8192	7.62E+05			122.37	
10		94	8192	7.70E+05			122.09	
11		96	8192	7.86E+05			122.60	
12		97	8192	7.95E+05			122.78	
13		99	8192	8.11E+05			123.10	
14		101	8192	8.27E+05			123.07	
15		102	8192	8.36E+05			123.18	
16		104	8192	8.52E+05			123.35	
17		105	8192	8.60E+05			123.30	
18		107	8192	8.77E+05			123.72	
19		109	8192	8.93E+05			123.50	
20		110	8192	9.01E+05			116.54	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
21		112	8192	9.18E+05			114.65	
22		113	8192	9.26E+05			115.65	
23		115	8192	9.42E+05			113.60	
24		116	8192	9.50E+05			115.91	
25		118	8192	9.67E+05			114.76	
26		120	8192	9.83E+05			115.66	
27		121	8192	9.91E+05			124.50	
28		123	8192	1.01E+06			123.97	
29		124	8192	1.02E+06			124.86	
30		126	8192	1.03E+06			125.29	
31		127	8192	1.04E+06			125.08	
32		129	8192	1.06E+06			125.03	
33		131	8192	1.07E+06			122.17	
34		132	8192	1.08E+06			120.15	
35		134	8192	1.10E+06			122.83	
36		135	8192	1.11E+06			123.50	
37		137	8192	1.12E+06			122.28	
38		138	8192	1.13E+06			120.44	
39		140	8192	1.15E+06			118.13	
40		142	8192	1.16E+06			112.98	
41		143	8192	1.17E+06			114.50	
42		145	8192	1.19E+06			114.93	
43		146	8192	1.20E+06			117.58	
44		148	8192	1.21E+06			117.42	
45		149	8192	1.22E+06			116.41	
46		151	8192	1.24E+06			115.61	
47		153	8192	1.25E+06			117.27	
48		154	8192	1.26E+06			116.51	
49		156	8192	1.28E+06			117.28	
50		157	8192	1.29E+06			115.65	
51		159	8192	1.30E+06			118.00	
52		160	8192	1.31E+06			123.63	
53		162	8192	1.33E+06			126.82	
54		164	8192	1.34E+06			126.67	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled	Disabled	
0		79	4096	3.24E+05			58.35	
1		80	4096	3.28E+05			60.99	
2		82	4096	3.36E+05			63.52	
3		83	4096	3.40E+05			63.42	
4		85	4096	3.48E+05			63.43	
5		86	4096	3.52E+05			63.34	
6		88	4096	3.60E+05			63.57	
7		90	4096	3.69E+05			63.77	
8		91	4096	3.73E+05			62.42	
9		93	4096	3.81E+05			62.48	
10		94	4096	3.85E+05			62.63	
11		96	4096	3.93E+05			62.00	
12		97	4096	3.97E+05			60.68	
13		99	4096	4.06E+05			58.58	
14		101	4096	4.14E+05			58.68	
15		102	4096	4.18E+05			59.37	
16		104	4096	4.26E+05			59.92	
17		105	4096	4.30E+05			60.03	
18		107	4096	4.38E+05			59.64	
19		109	4096	4.46E+05			58.98	
20		110	4096	4.51E+05			58.38	
21		112	4096	4.59E+05			58.92	
22		113	4096	4.63E+05			59.27	
23		115	4096	4.71E+05			58.99	
24		116	4096	4.75E+05			59.58	
25		118	4096	4.83E+05			63.17	
26		120	4096	4.92E+05			60.11	
27		121	4096	4.96E+05			62.97	
28		123	4096	5.04E+05			64.25	
29		124	4096	5.08E+05			64.46	
30		126	4096	5.16E+05			64.55	
31		127	4096	5.20E+05			64.46	
32		129	4096	5.28E+05			64.54	
33		131	4096	5.37E+05			64.53	
34		132	4096	5.41E+05			64.24	
35		134	4096	5.49E+05			64.60	
36		135	4096	5.53E+05			64.68	
37		137	4096	5.61E+05			64.49	
38		138	4096	5.65E+05			64.55	
39		140	4096	5.73E+05			64.64	
40		142	4096	5.82E+05			64.64	
41		143	4096	5.86E+05			64.47	
42		145	4096	5.94E+05			64.69	
43		146	4096	5.98E+05			63.54	
44		148	4096	6.06E+05			63.49	
45		149	4096	6.10E+05			63.50	
46		151	4096	6.18E+05			63.51	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
47		153	4096	6.27E+05			63.52	
48		154	4096	6.31E+05			63.66	
49		156	4096	6.39E+05			63.90	
50		157	4096	6.43E+05			64.21	
51		159	4096	6.51E+05			63.67	
52		160	4096	6.55E+05			62.46	
53		162	4096	6.64E+05			62.29	
54		164	4096	6.72E+05			62.11	


Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		79	2048	1.62E+05			31.57	
1		80	2048	1.64E+05			31.30	
2		82	2048	1.68E+05			31.75	
3		83	2048	1.70E+05			31.07	
4		85	2048	1.74E+05			30.77	
5		86	2048	1.76E+05			31.07	
6		88	2048	1.80E+05			30.54	
7		90	2048	1.84E+05			30.61	
8		91	2048	1.86E+05			30.54	
9		93	2048	1.90E+05			30.42	
10		94	2048	1.93E+05			30.61	
11		96	2048	1.97E+05			30.33	
12		97	2048	1.99E+05			30.15	
13		99	2048	2.03E+05			30.20	
14		101	2048	2.07E+05			29.53	
15		102	2048	2.09E+05			29.04	
16		104	2048	2.13E+05			29.24	
17		105	2048	2.15E+05			29.14	
18		107	2048	2.19E+05			29.30	
19		109	2048	2.23E+05			29.91	
20		110	2048	2.25E+05			30.00	
21		112	2048	2.29E+05			30.15	
22		113	2048	2.31E+05			30.22	
23		115	2048	2.36E+05			30.21	
24		116	2048	2.38E+05			30.26	
25		118	2048	2.42E+05			30.27	
26		120	2048	2.46E+05			30.21	
27		121	2048	2.48E+05			30.29	
28		123	2048	2.52E+05			30.52	
29		124	2048	2.54E+05			30.82	
30		126	2048	2.58E+05			31.06	
31		127	2048	2.60E+05			31.57	
32		129	2048	2.64E+05			31.79	
33		131	2048	2.68E+05			32.07	


Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
34		132	2048	2.70E+05			32.17	
35		134	2048	2.74E+05			32.26	
36		135	2048	2.76E+05			32.60	
37		137	2048	2.81E+05			32.64	
38		138	2048	2.83E+05			32.70	
39		140	2048	2.87E+05			32.77	
40		142	2048	2.91E+05			32.66	
41		143	2048	2.93E+05			32.83	
42		145	2048	2.97E+05			32.80	
43		146	2048	2.99E+05			32.82	
44		148	2048	3.03E+05			32.83	
45		149	2048	3.05E+05			32.78	
46		151	2048	3.09E+05			32.84	
47		153	2048	3.13E+05			32.84	
48		154	2048	3.15E+05			32.86	
49		156	2048	3.19E+05			32.86	
50		157	2048	3.22E+05			32.85	
51		159	2048	3.26E+05			32.84	
52		160	2048	3.28E+05			32.69	
53		162	2048	3.32E+05			32.70	
54		164	2048	3.36E+05			32.85	


Hitachi OD152S-1 Average Read Rate




Hitachi OD152S-1 Average Write Test




Write Test 128k Block (No Cache)




Write Test 64k Block (No Cache)




Write Test 32k Block (No Cache)




Write Test 16k Block (No Cache)




Write Test 8k Block (No Cache)




Write Test 4k Block (No Cache)




Write Test 2k Block (No Cache)




Read Cache Test (128k Block)





Read Cache Test (64k Block)




Read Cache Test (32k Block)






Read Cache Test 4k Block



Read Cache Test (2k Block)



Optical Drive Description:

Maxoptics TMT 3-1301 Standard, 5.25 inch, full height

Interface: Single Ended Cache: 1 M Rotational Speed: 3375

Power Requirements:

	+5 VDC +/- 5%	+12 VDC +/- 5%
Tolerance		
Ripple	50 mV PP (max)	120 mV PP (max)
Current: Typical	2.0 A	1.0 A
Maximum	2.2 A	4.0 A

2 Pass Write: Erase/Write

3 Pass Write: Erase/Write/Verify

Media Description:

Maxell MA-132-S1 Format: ZCAV Type: 512 Byte/sector Formatted Capacity: 1.19 GB

Read Test:

Read Cache Enabled

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	1169.60
9088	64	581632	981.52
18176	32	581632	836.94
36531	16	584496	468.95
72702	8	581616	343.90
145404	4	581616	301.59
290808	2	581616	267.45

Read Cache Disabled

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	1169.34
9088	64	581632	982.16
18176	32	581632	835.23
36531	16	584496	469.56
72702	8	581616	344.45
145404	4	581616	303.22
290808	2	581616	268.10

Write Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache			X	
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	607.44
9088	64	581632	521.01
18176	32	581632	551.79
36531	16	584496	501.25
72702	8	581616	184.62
145404	4	581616	98.08
290808	2	581616	80.30

Write Test 2:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache			X	
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	606.67
9088	64	581632	522.32
18176	32	581632	552.03
36531	16	584496	500.56
72702	8	581616	184.87
145404	4	581616	98.03
290808	2	581616	80.41

Write Cache Test 1:

Conditions	3 Pass	2 Pass	Read	Write
Write Mode		X		
Test Mode				X
Burst Size: unlimited				X

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Cache Disabled	Cache Enabled
4	131072	524288	2327.27	3100.76
8	131072	1048576	2089.80	2089.80
12	131072	1572864	1269.42	1462.86
8	65536	524288	2327.27	2327.27
16	65536	1048576	1861.82	2089.80
24	65536	1572864	1121.17	1074.13
16	32768	524288	2327.27	2327.27
32	32768	1048576	1163.64	1551.52
48	32768	1572864	1121.17	1074.13
32	16384	524288	721.13	1312.82
64	16384	1048576	533.33	533.33
96	16384	1572864	527.84	527.84
64	8192	524288	266.67	275.27
128	8192	1048576	252.22	251.60
192	8192	1572864	225.55	223.91
128	4096	524288	122.78	121.04
256	4096	1048576	120.33	115.84
384	4096	1572864	112.78	110.11
256	2048	524288	290.91	246.15
512	2048	1048576	211.57	184.50
768	2048	1572864	194.18	142.62

Block Write Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache			X	X

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		194	131072	2.54E+07	477.45		476.99
1		206	131072	2.70E+07	462.03		461.54
2		218	131072	2.86E+07	494.23		494.23
3		231	131072	3.03E+07	495.69		502.17
4		243	131072	3.19E+07	508.82		512.93
5		255	131072	3.34E+07	563.24		560.15
6		267	131072	3.50E+07	582.02		568.75
7		279	131072	3.66E+07	583.15		585.83
8		291	131072	3.81E+07	594.92		599.04
9		303	131072	3.97E+07	684.87		684.14
10		315	131072	4.13E+07	699.15		697.22
11		327	131072	4.29E+07	676.19		686.50
12		340	131072	4.46E+07	727.64		728.25
13		352	131072	4.61E+07	698.76		704.77
14		364	131072	4.77E+07	736.98		738.27
15		367	131072	4.81E+07	713.92		713.27

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		388	65536	2.54E+07	441.93		441.07
1		412	65536	2.70E+07	432.12		432.12
2		436	65536	2.86E+07	422.66		424.07
3		461	65536	3.02E+07	445.07		447.23
4		485	65536	3.18E+07	460.19		459.44
5		509	65536	3.34E+07	448.95		447.23
6		533	65536	3.49E+07	481.06		481.06
7		558	65536	3.66E+07	531.19		527.82
8		582	65536	3.81E+07	543.85		543.85
9		606	65536	3.97E+07	576.89		577.40
10		630	65536	4.13E+07	544.94		547.38
11		654	65536	4.29E+07	588.94		587.04
12		679	65536	4.45E+07	596.67		603.05
13		703	65536	4.61E+07	564.52		568.44
14		727	65536	4.76E+07	573.15		566.65
15		733	65536	4.80E+07	620.23		620.64

Block Write Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)	
						Enabled	Disabled
0		775	32768	2.54E+07		435.39	448.38
1		824	32768	2.70E+07		442.86	444.13
2		872	32768	2.86E+07		466.15	455.20
3		921	32768	3.02E+07		501.05	488.68
4		969	32768	3.18E+07		513.21	508.16
5		1018	32768	3.34E+07		521.63	509.16
6		1066	32768	3.49E+07		540.95	541.89
7		1115	32768	3.65E+07		570.88	543.57
8		1163	32768	3.81E+07		583.14	565.59
9		1211	32768	3.97E+07		585.55	578.73
10		1260	32768	4.13E+07		599.29	581.23
11		1308	32768	4.29E+07		613.54	590.77
12		1357	32768	4.45E+07		624.00	597.55
13		1405	32768	4.60E+07		617.75	594.00
14		1454	32768	4.76E+07		620.13	605.52
15		1465	32768	4.80E+07		643.60	629.87

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)	
						Enabled	Disabled
0		1550	16384	2.54E+07		401.75	388.90
1		1647	16384	2.70E+07		405.17	399.15
2		1744	16384	2.86E+07		400.69	393.51
3		1841	16384	3.02E+07		419.00	410.99
4		1938	16384	3.18E+07		425.41	424.48
5		2035	16384	3.33E+07		457.75	452.47
6		2132	16384	3.49E+07		465.57	196.16
7		2229	16384	3.65E+07		466.14	447.20
8		2325	16384	3.81E+07		496.93	488.70
9		2422	16384	3.97E+07		505.77	490.97
10		2519	16384	4.13E+07		524.11	472.50
11		2616	16384	4.29E+07		537.79	501.03
12		2713	16384	4.44E+07		261.53	546.91
13		2810	16384	4.60E+07		534.28	265.34
14		2907	16384	4.76E+07		563.37	540.77
15		2930	16384	4.80E+07		552.61	536.38

Block Write Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		3100	8192	2.54E+07		153.74		165.45
1		3294	8192	2.70E+07		176.33		174.46
2		3488	8192	2.86E+07		175.18		183.47
3		3682	8192	3.02E+07		192.64		193.40
4		3875	8192	3.17E+07		162.56		157.92
5		4096	8192	3.36E+07		198.94		199.28
6		4263	8192	3.49E+07		175.74		179.77
7		4457	8192	3.65E+07		200.43		199.06
8		4650	8192	3.81E+07		175.10		176.56
9		4844	8192	3.97E+07		199.59		197.02
10		5038	8192	4.13E+07		181.81		180.03
11		5232	8192	4.29E+07		179.43		180.66
12		5425	8192	4.44E+07		202.29		213.33
13		5619	8192	4.60E+07		184.04		182.15
14		5813	8192	4.76E+07		195.08		205.40
15		5859	8192	4.80E+07		188.00		186.40

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		6200	4096	2.54E+07		107.94		107.76
1		6588	4096	2.70E+07		107.67		110.07
2		6975	4096	2.86E+07		88.66		94.19
3		7363	4096	3.02E+07		88.08		88.43
4		7750	4096	3.17E+07		84.19		86.26
5		8138	4096	3.33E+07		104.49		105.48
6		8525	4096	3.49E+07		107.97		104.19
7		8913	4096	3.65E+07		103.20		102.64
8		9300	4096	3.81E+07		96.58		96.44
9		9688	4096	3.97E+07		101.44		107.16
10		10075	4096	4.13E+07		107.13		109.45
11		10463	4096	4.29E+07		101.65		102.98
12		10850	4096	4.44E+07		89.60		89.90
13		11238	4096	4.60E+07		112.31		111.67
14		11625	4096	4.76E+07		93.99		87.99
15		11717	4096	4.80E+07		104.44		103.73

Block Write Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		12400	2048	2.54E+07		104.25		96.75
1		13175	2048	2.70E+07		62.55		44.97
2		13950	2048	2.86E+07		117.26		137.85
3		14725	2048	3.02E+07		120.25		128.74
4		15500	2048	3.17E+07		82.50		77.30
5		16275	2048	3.33E+07		46.46		44.97
6		17050	2048	3.49E+07		63.68		63.05
7		17825	2048	3.65E+07		117.33		122.42
8		18600	2048	3.81E+07		137.55		143.61
9		19375	2048	3.97E+07		72.74		70.05
10		20150	2048	4.13E+07		47.77		48.47
11		20925	2048	4.29E+07		85.92		81.35
12		21700	2048	4.44E+07		84.05		87.43
13		22475	2048	4.60E+07		52.35		62.01
14		23250	2048	4.76E+07		120.69		130.53
15		23433	2048	4.80E+07		122.61		122.27

Write Test 3:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode	X			
Write Cache			X	
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	407.86
9088	64	581632	350.80
18176	32	581632	400.87
36531	16	584496	358.45
72702	8	581616	134.45
145404	4	581616	70.29
290808	2	581616	58.96

Write Test 4:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode	X			
Write Cache				X
SCSI Parity				X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	408.37
9088	64	581632	351.07
18176	32	581632	400.58
36531	16	584496	357.47
72702	8	581616	134.33
145404	4	581616	70.89
290808	2	581616	58.92

Write Cache Test 2:

Conditions	3 Pass	2 Pass	Read	Write
Write Mode	X			
Test Mode				X
Burst Size: unlimited				X

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Cache Disabled	Cache Enabled
4	131072	524288	3011.76	3011.76
8	131072	1048576	1706.67	1706.67
12	131072	1572864	966.04	843.96
8	65536	524288	2327.27	3200.00
16	65536	1048576	1861.82	1861.82
24	65536	1572864	872.73	898.25
16	32768	524288	2327.27	2327.27
32	32768	1048576	1442.25	1551.52
48	32768	1572864	930.91	848.62
32	16384	524288	930.91	492.31
64	16384	1048576	372.36	372.36
96	16384	1572864	307.82	310.93
64	8192	524288	216.95	211.57
128	8192	1048576	200.39	204.80
192	8192	1572864	178.19	179.44
128	4096	524288	75.07	75.18
256	4096	1048576	72.01	72.27
384	4096	1572864	69.38	69.22
256	2048	524288	79.01	119.63
512	2048	1048576	59.78	79.63
768	2048	1572864	54.62	55.05

Block Write Test 2:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode	X			
Write Cache			X	X

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		194	131072	2.54E+07		317.71		321.33
1		206	131072	2.70E+07		305.19		305.57
2		218	131072	2.86E+07		331.83		330.30
3		231	131072	3.03E+07		336.04		336.46
4		243	131072	3.19E+07		342.18		343.43
5		255	131072	3.34E+07		382.16		379.71
6		267	131072	3.50E+07		393.05		391.57
7		279	131072	3.66E+07		392.40		392.18
8		291	131072	3.81E+07		396.13		398.67
9		303	131072	3.97E+07		464.59		462.71
10		315	131072	4.13E+07		469.38		471.19
11		327	131072	4.29E+07		462.70		459.91
12		340	131072	4.46E+07		484.04		485.23
13		352	131072	4.61E+07		473.33		470.61
14		364	131072	4.77E+07		484.78		484.73
15		367	131072	4.81E+07		481.56		479.10

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		388	65536	2.54E+07		299.80		295.09
1		412	65536	2.70E+07		292.20		293.24
2		436	65536	2.86E+07		287.02		287.67
3		461	65536	3.02E+07		298.77		298.44
4		485	65536	3.18E+07		304.49		303.84
5		509	65536	3.34E+07		298.04		299.25
6		533	65536	3.49E+07		314.77		314.48
7		558	65536	3.66E+07		367.14		366.50
8		582	65536	3.81E+07		385.31		385.75
9		606	65536	3.97E+07		389.48		388.62
10		630	65536	4.13E+07		374.93		374.93
11		654	65536	4.29E+07		394.83		394.01
12		679	65536	4.45E+07		400.41		402.04
13		703	65536	4.61E+07		382.07		382.07
14		727	65536	4.76E+07		383.99		383.48
15		733	65536	4.80E+07		406.33		405.18

Block Write Test 2 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		775	32768	2.54E+07	306.97		316.65
1		824	32768	2.70E+07	311.94		311.31
2		872	32768	2.86E+07	325.91		329.25
3		921	32768	3.02E+07	362.78		358.41
4		969	32768	3.18E+07	372.87		378.10
5		1018	32768	3.34E+07	379.94		382.39
6		1066	32768	3.49E+07	390.12		394.81
7		1115	32768	3.65E+07	406.29		404.49
8		1163	32768	3.81E+07	422.14		425.86
9		1211	32768	3.97E+07	431.01		426.83
10		1260	32768	4.13E+07	431.55		432.06
11		1308	32768	4.29E+07	432.26		438.70
12		1357	32768	4.45E+07	449.71		442.65
13		1405	32768	4.60E+07	441.30		443.39
14		1454	32768	4.76E+07	447.73		450.15
15		1465	32768	4.80E+07	458.53		466.27

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		1550	16384	2.54E+07	286.14		282.20
1		1647	16384	2.70E+07	292.57		291.31
2		1744	16384	2.86E+07	278.37		284.13
3		1841	16384	3.02E+07	381.66		293.85
4		1938	16384	3.18E+07	303.35		317.35
5		2035	16384	3.33E+07	323.59		318.72
6		2132	16384	3.49E+07	339.56		326.03
7		2229	16384	3.65E+07	319.40		343.39
8		2325	16384	3.81E+07	375.23		365.53
9		2422	16384	3.97E+07	379.96		369.98
10		2519	16384	4.13E+07	383.59		376.11
11		2616	16384	4.29E+07	213.57		395.88
12		2713	16384	4.44E+07	409.08		400.37
13		2810	16384	4.60E+07	405.23		402.04
14		2907	16384	4.76E+07	430.95		415.73
15		2930	16384	4.80E+07	298.75		231.10

Block Write Test 2 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		3100	8192	2.54E+07		134.90		136.08
1		3294	8192	2.70E+07		125.80		125.70
2		3488	8192	2.86E+07		116.36		115.67
3		3682	8192	3.02E+07		134.74		133.67
4		3875	8192	3.17E+07		132.83		132.39
5		4096	8192	3.36E+07		136.02		136.18
6		4263	8192	3.49E+07		139.75		139.75
7		4457	8192	3.65E+07		118.72		118.46
8		4650	8192	3.81E+07		139.56		138.42
9		4844	8192	3.97E+07		127.38		127.28
10		5038	8192	4.13E+07		142.79		140.47
11		5232	8192	4.29E+07		139.16		138.35
12		5425	8192	4.44E+07		128.84		128.96
13		5619	8192	4.60E+07		129.44		129.40
14		5813	8192	4.76E+07		129.26		129.62
15		5859	8192	4.80E+07		160.15		160.30



Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		6200	4096	2.54E+07		74.12		73.55
1		6588	4096	2.70E+07		73.23		73.63
2		6975	4096	2.86E+07		67.54		67.28
3		7363	4096	3.02E+07		77.01		76.06
4		7750	4096	3.17E+07		69.47		69.51
5		8138	4096	3.33E+07		62.83		62.77
6		8525	4096	3.49E+07		71.29		71.29
7		8913	4096	3.65E+07		70.15		71.46
8		9300	4096	3.81E+07		61.86		66.03
9		9688	4096	3.97E+07		81.26		81.31
10		10075	4096	4.13E+07		66.68		70.24
11		10463	4096	4.29E+07		68.19		68.59
12		10850	4096	4.44E+07		80.09		79.64
13		11238	4096	4.60E+07		71.20		71.20
14		11625	4096	4.76E+07		70.80		70.66
15		11717	4096	4.80E+07		62.26		62.26


Block Write Test 2 (Continued)

Block Transfer Rate


Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		12400	2048	2.54E+07		32.63		32.43
1		13175	2048	2.70E+07		87.77		83.42
2		13950	2048	2.86E+07		44.21		46.80
3		14725	2048	3.02E+07		58.99		60.41
4		15500	2048	3.17E+07		70.23		69.13
5		16275	2048	3.33E+07		80.54		78.45
6		17050	2048	3.49E+07		72.98		77.54
7		17825	2048	3.65E+07		38.41		35.77
8		18600	2048	3.81E+07		32.58		32.85
9		19375	2048	3.97E+07		51.07		52.21
10		20150	2048	4.13E+07		96.70		97.12
11		20925	2048	4.29E+07		74.06		74.08
12		21700	2048	4.44E+07		99.33		99.58
13		22475	2048	4.60E+07		57.92		55.57
14		23250	2048	4.76E+07		46.56		46.13
15		23433	2048	4.80E+07		63.41		64.42




Maxoptics TMT3-1301 2-Pass Write Cache, 64k Block




Maxoptics TMT3-1301 2-Pass Write Cache, 32k Block




Maxoptics TMT3-1301 2-Pass Write Cache, 16k Block




Maxoptics TMT3-1301 2-Pass Write Cache, 8k Block




Maxoptics TMT3-1301 2-Pass Write Cache, 4k Block




Maxoptics TMT3-1301 2-Pass Write Cache, 2k Block




Maxoptics TMT3-1301 2-Pass Band Write (Cache Disabled)




Maxoptics TMT3-1301 2-Pass Band Write (Cache Enabled)





Maxoptics TMT3-1301 Cache Effects, 128k Block 2-Pass Write




Maxoptics TMT3-1301 Cache Effects, 64k Block 2-Pass Write







Maxoptics TMT3-1301 Cache Effects, 8k Block 2-Pass Write



Maxoptics TMT3-1301 Cache Effects, 4k Block 2-Pass Write



Maxoptics TMT3-1301 Cache Effects, 2k Block 2-Pass Write



Read Cache Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode				
Read Cache			X	X

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
				Enabled		Disabled
0	194	131072	2.54E+07	883.07		914.96
1	206	131072	2.70E+07	889.01		867.94
2	218	131072	2.86E+07	940.80		939.21
3	231	131072	3.03E+07	978.75		976.81
4	243	131072	3.19E+07	1044.81		1043.06
5	255	131072	3.34E+07	1051.89		1051.55
6	267	131072	3.50E+07	1108.89		1103.52
7	279	131072	3.66E+07	1134.80		1128.70
8	291	131072	3.81E+07	1173.17		1173.17
9	303	131072	3.97E+07	1251.90		1245.47
10	315	131072	4.13E+07	1287.77		1332.45
11	327	131072	4.29E+07	1285.11		1257.69
12	340	131072	4.46E+07	1456.49		1456.49
13	352	131072	4.61E+07	1380.82		1380.82
14	364	131072	4.77E+07	1386.25		1388.32
15	367	131072	4.81E+07	1448.33		1448.33

Block Transfer Rate

Band Number	Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
				Enabled		Disabled
0	388	65536	2.54E+07	798.71		798.97
1	412	65536	2.70E+07	837.88		837.88
2	436	65536	2.86E+07	811.40		803.92
3	461	65536	3.02E+07	793.33		813.68
4	485	65536	3.18E+07	946.63		927.95
5	509	65536	3.34E+07	868.46		869.62
6	533	65536	3.49E+07	1023.16		1023.16
7	558	65536	3.66E+07	917.10		908.01
8	582	65536	3.81E+07	957.78		957.78
9	606	65536	3.97E+07	1115.44		1115.44
10	630	65536	4.13E+07	1005.49		1009.77
11	654	65536	4.29E+07	996.10		1018.64
12	679	65536	4.45E+07	1208.12		1189.92
13	703	65536	4.61E+07	1079.20		1067.93
14	727	65536	4.76E+07	1086.09		1088.88
15	733	65536	4.80E+07	1227.37		1222.07

Read Cache Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		775	32768	2.54E+07	656.26		652.46
1		824	32768	2.70E+07	808.09		812.32
2		872	32768	2.86E+07	734.12		747.10
3		921	32768	3.02E+07	816.85		844.95
4		969	32768	3.18E+07	780.86		766.01
5		1018	32768	3.34E+07	725.85		702.67
6		1066	32768	3.49E+07	867.33		863.81
7		1115	32768	3.65E+07	731.60		729.95
8		1163	32768	3.81E+07	863.08		851.24
9		1211	32768	3.97E+07	949.57		945.86
10		1260	32768	4.13E+07	746.81		748.33
11		1308	32768	4.29E+07	1011.99		1003.98
12		1357	32768	4.45E+07	931.25		926.87
13		1405	32768	4.60E+07	956.19		957.41
14		1454	32768	4.76E+07	757.79		755.57
15		1465	32768	4.80E+07	972.37		970.21

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		1550	16384	2.54E+07	461.74		464.51
1		1647	16384	2.70E+07	371.89		334.12
2		1744	16384	2.86E+07	398.46		398.74
3		1841	16384	3.02E+07	498.83		506.46
4		1938	16384	3.18E+07	406.18		409.40
5		2035	16384	3.33E+07	535.97		533.60
6		2132	16384	3.49E+07	492.52		502.53
7		2229	16384	3.65E+07	425.53		426.09
8		2325	16384	3.81E+07	431.40		433.62
9		2422	16384	3.97E+07	449.66		445.68
10		2519	16384	4.13E+07	632.02		633.11
11		2616	16384	4.29E+07	452.79		451.72
12		2713	16384	4.44E+07	576.85		578.16
13		2810	16384	4.60E+07	523.70		512.60
14		2907	16384	4.76E+07	475.49		474.13
15		2930	16384	4.80E+07	470.71		475.32

Read Cache Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		3100	8192	2.54E+07		331.99		325.08
1		3294	8192	2.70E+07		270.14		262.31
2		3488	8192	2.86E+07		279.29		272.69
3		3682	8192	3.02E+07		284.79		286.79
4		3875	8192	3.17E+07		333.55		332.40
5		4096	8192	3.36E+07		356.34		356.15
6		4263	8192	3.49E+07		432.08		308.61
7		4457	8192	3.65E+07		346.24		346.04
8		4650	8192	3.81E+07		350.38		352.04
9		4844	8192	3.97E+07		370.76		351.91
10		5038	8192	4.13E+07		393.25		390.13
11		5232	8192	4.29E+07		346.55		323.44
12		5425	8192	4.44E+07		360.47		364.12
13		5619	8192	4.60E+07		403.95		403.34
14		5813	8192	4.76E+07		459.66		450.84
15		5859	8192	4.80E+07		374.56		374.38

Block Transfer Rate


Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		6200	4096	2.54E+07		265.75		264.99
1		6588	4096	2.70E+07		276.20		274.96
2		6975	4096	2.86E+07		277.86		271.64
3		7363	4096	3.02E+07		279.43		278.85
4		7750	4096	3.17E+07		303.30		296.74
5		8138	4096	3.33E+07		304.08		306.29
6		8525	4096	3.49E+07		322.00		323.50
7		8913	4096	3.65E+07		290.82		290.18
8		9300	4096	3.81E+07		292.68		292.04
9		9688	4096	3.97E+07		261.70		288.20
10		10075	4096	4.13E+07		290.68		283.84
11		10463	4096	4.29E+07		347.78		320.83
12		10850	4096	4.44E+07		299.31		300.22
13		11238	4096	4.60E+07		314.28		314.79
14		11625	4096	4.76E+07		335.55		337.15
15		11717	4096	4.80E+07		329.26		326.33

Read Cache Test 1 (Continued)


Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Transfer Rate (Cache)		
					Enabled		Disabled
0		12400	2048	2.54E+07	242.88		239.66
1		13175	2048	2.70E+07	311.72		305.37
2		13950	2048	2.86E+07	251.08		241.77
3		14725	2048	3.02E+07	317.45		251.39
4		15500	2048	3.17E+07	201.50		204.04
5		16275	2048	3.33E+07	266.34		248.28
6		17050	2048	3.49E+07	270.03		270.76
7		17825	2048	3.65E+07	269.00		267.99
8		18600	2048	3.81E+07	275.09		269.94
9		19375	2048	3.97E+07	271.66		271.13
10		20150	2048	4.13E+07	245.13		271.14
11		20925	2048	4.29E+07	275.87		196.12
12		21700	2048	4.44E+07	303.92		308.19
13		22475	2048	4.60E+07	288.36		284.76
14		23250	2048	4.76E+07	279.23		280.42
15		23433	2048	4.80E+07	280.93		283.85


Maxoptics TMT3-1301 Average Write Rate (3 Pass)




Maxoptics MTM3-1301 3-Pass Write Cache, 128k Block




Maxoptics MTM3-1301 3-Pass Write Cache, 64k Block




Maxoptics MTM3-1301 3-Pass Write Cache, 32k Block




Maxoptics MTM3-1301 3-Pass Write Cache, 16k Block




Maxoptics MTM3-1301 3-Pass Write Cache, 8k Block




Maxoptics MTM3-1301 3-Pass Write Cache, 4k Block




Maxoptics MTM3-1301 3-Pass Write Cache, 2k Block




Maxoptics TMT3-1301 3-Pass Band Write (Cache Disabled)




Maxoptics TMT3-1301 3-Pass Band Write (Cache Enabled)




Maxoptics TMT3-1301 Cache Effects, 128k Block 3-Pass Write




Maxoptics TMT3-1301 Cache Effects, 64k Block 3-Pass Write




Maxoptics TMT3-1301 Cache Effects, 32k Block 3-Pass Write




Maxoptics TMT3-1301 Cache Effects, 16k Block 3-Pass Write




Maxoptics TMT3-1301 Cache Effects, 8k Block 3-Pass Write




Maxoptics TMT3-1301 Cache Effects, 4k Block 3-Pass Write




Maxoptics TMT3-1301 Cache Effects, 2k Block 3-Pass Write




Maxoptics TMT3-1301 Average Read Test



Maxoptics TMT3-1301 Band Read (Cache Disabled)



Maxoptics TMT3-1301 Band Read (Cache Enabled)



Optical Drive Description:

Pinnacle Sierra Standard, 5.25 inch, full height

Interface: Single Ended Cache: 1 M Rotational Speed: 3000

Power Requirements:

+5 VDC +12 VDC

Tolerance
Ripple

Current: Typical
Maximum

2 Pass Write: Erase/Write

Media Description:

Maxell MA-132-S1 Format: ZCAV Type: 512 Byte/sector Formatted Capacity: 1.19 GB

Read Test 1:

Conditions	Enabled	Disabled
Read Cache	X	
Disconnect		X
SCSI Parity		X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	1026.83
9088	64	581632	886.33
18176	32	581632	672.25
36531	16	584496	489.46
72702	8	581616	289.18
145404	4	581616	167.78
290808	2	581616	92.87

Read Test 2:

Conditions	Enabled	Disabled
Read Cache	X	
Disconnect	X	
SCSI Parity		X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	1000.36
9088	64	581632	881.89
18176	32	581632	704.07
36531	16	584496	514.62
72702	8	581616	312.55
145404	4	581616	169.71
290808	2	581616	91.13

Write Test 1:

Conditions	Enabled	Disabled
Write Cache	X	
Disconnect		X
SCSI Parity		X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	539.08
9088	64	581632	488.39
18176	32	581632	481.26
36531	16	584496	441.23
72702	8	581616	391.06
145404	4	581616	198.35
290808	2	581616	184.82

Write Test 2:

Conditions	Enabled	Disabled
Write Cache	X	
Disconnect	X	
SCSI Parity		X

Average Transfer Rate

Number of CDBs	Block Size (kB)	Total Stored	Transfer Rate
4544	128	581632	539.49
9088	64	581632	505.11
18176	32	581632	434.11
36531	16	584496	423.53
72702	8	581616	360.10
145404	4	581616	223.77
290808	2	581616	131.80

Write Cache Test 1:

Conditions	2 Pass	Enabled	Disabled
Write Mode	X		
Disconnect			
Cache		X	

Average Transfer Rate

Number of CDBs	Block Size	Total Stored	Connect Disabled	Connect Enabled
4	131072	524288	1878.90	1692.56
8	131072	1048576	2133.33	1735.59
12	131072	1572864	955.52	1107.03
8	65536	524288	1374.50	1374.50
16	65536	1048576	2017.73	1625.40
24	65536	1572864	972.15	917.01
16	32768	524288	1780.87	1163.64
32	32768	1048576	1773.16	1462.86
48	32768	1572864	741.13	716.92
32	16384	524288	1204.71	1131.49
64	16384	1048576	1128.37	1163.64
96	16384	1572864	654.31	589.07
64	8192	524288	1131.49	906.19
128	8192	1048576	920.45	819.20
192	8192	1572864	579.62	375.55
128	4096	524288	809.49	664.94
256	4096	1048576	463.35	649.13
384	4096	1572864	201.57	311.72
256	2048	524288	455.11	369.68
512	2048	1048576	378.56	356.79
768	2048	1572864	270.18	171.57

Block Write Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Write Mode		X		
Write Cache			X	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Disconnect	
						Enabled	Disabled
0		194	131072	2.54E+07		506.26	471.91
1		206	131072	2.70E+07		424.47	432.55
2		218	131072	2.86E+07		433.43	459.33
3		231	131072	3.03E+07		466.52	486.80
4		243	131072	3.19E+07		479.11	486.53
5		255	131072	3.34E+07		471.27	510.08
6		267	131072	3.50E+07		513.77	522.41
7		279	131072	3.66E+07		529.93	544.97
8		291	131072	3.81E+07		531.89	554.95
9		303	131072	3.97E+07		558.20	586.48
10		315	131072	4.13E+07		556.91	591.03
11		327	131072	4.29E+07		574.24	583.93
12		340	131072	4.46E+07		584.32	621.89
13		352	131072	4.61E+07		589.28	651.00
14		364	131072	4.77E+07		606.82	643.09
15		367	131072	4.81E+07		621.13	664.07

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Disconnect	
						Enabled	Disabled
0		388	65536	2.54E+07		458.58	426.08
1		412	65536	2.70E+07		404.42	392.85
2		436	65536	2.86E+07		421.96	425.50
3		461	65536	3.02E+07		429.40	459.13
4		485	65536	3.18E+07		453.20	476.15
5		509	65536	3.34E+07		459.79	468.79
6		533	65536	3.49E+07		484.82	520.95
7		558	65536	3.66E+07		476.35	521.04
8		582	65536	3.81E+07		503.42	552.64
9		606	65536	3.97E+07		518.09	555.09
10		630	65536	4.13E+07		526.92	583.50
11		654	65536	4.29E+07		538.96	572.98
12		679	65536	4.45E+07		550.56	583.93
13		703	65536	4.61E+07		568.51	607.26
14		727	65536	4.76E+07		563.57	600.75
15		733	65536	4.80E+07		584.21	615.27

Block Write Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Disconnect	
						Enabled	Disabled
0		775	32768	2.54E+07		397.12	414.99
1		824	32768	2.70E+07		363.10	392.85
2		872	32768	2.86E+07		373.55	382.51
3		921	32768	3.02E+07		381.07	404.06
4		969	32768	3.18E+07		393.40	411.74
5		1018	32768	3.34E+07		408.48	451.32
6		1066	32768	3.49E+07		435.82	448.13
7		1115	32768	3.65E+07		421.30	448.64
8		1163	32768	3.81E+07		438.25	467.95
9		1211	32768	3.97E+07		451.97	481.93
10		1260	32768	4.13E+07		451.76	497.35
11		1308	32768	4.29E+07		477.48	503.99
12		1357	32768	4.45E+07		482.38	537.09
13		1405	32768	4.60E+07		488.38	533.59
14		1454	32768	4.76E+07		489.36	543.04
15		1465	32768	4.80E+07		509.65	557.13

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Disconnect	
						Enabled	Disabled
0		1550	16384	2.54E+07		357.50	383.25
1		1647	16384	2.70E+07		335.95	356.98
2		1744	16384	2.86E+07		345.86	354.07
3		1841	16384	3.02E+07		365.60	371.17
4		1938	16384	3.18E+07		342.97	380.70
5		2035	16384	3.33E+07		363.03	403.52
6		2132	16384	3.49E+07		343.14	396.61
7		2229	16384	3.65E+07		355.43	420.27
8		2325	16384	3.81E+07		395.83	432.81
9		2422	16384	3.97E+07		411.16	453.72
10		2519	16384	4.13E+07		372.29	454.69
11		2616	16384	4.29E+07		361.67	443.81
12		2713	16384	4.44E+07		420.82	486.64
13		2810	16384	4.60E+07		416.99	473.96
14		2907	16384	4.76E+07		443.82	502.24
15		2930	16384	4.80E+07		462.12	521.31

Block Write Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Disconnect	
						Enabled	Disabled
0		3100	8192	2.54E+07		256.54	313.33
1		3294	8192	2.70E+07		267.29	292.93
2		3488	8192	2.86E+07		242.16	329.68
3		3682	8192	3.02E+07		259.82	336.87
4		3875	8192	3.17E+07		259.26	356.32
5		4096	8192	3.36E+07		269.00	369.28
6		4263	8192	3.49E+07		288.94	376.96
7		4457	8192	3.65E+07		268.35	387.10
8		4650	8192	3.81E+07		352.74	398.37
9		4844	8192	3.97E+07		381.79	418.26
10		5038	8192	4.13E+07		294.23	415.03
11		5232	8192	4.29E+07		319.63	416.23
12		5425	8192	4.44E+07		294.28	436.58
13		5619	8192	4.60E+07		296.87	437.19
14		5813	8192	4.76E+07		317.22	445.14
15		5859	8192	4.80E+07		304.97	462.54

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Disconnect	
						Enabled	Disabled
0		6200	4096	2.54E+07		280.10	308.42
1		6588	4096	2.70E+07		222.94	202.52
2		6975	4096	2.86E+07		210.17	169.72
3		7363	4096	3.02E+07		219.14	181.82
4		7750	4096	3.17E+07		216.90	200.57
5		8138	4096	3.33E+07		214.34	207.15
6		8525	4096	3.49E+07		219.46	193.48
7		8913	4096	3.65E+07		226.88	192.32
8		9300	4096	3.81E+07		243.46	187.72
9		9688	4096	3.97E+07		250.55	186.31
10		10075	4096	4.13E+07		242.55	188.95
11		10463	4096	4.29E+07		245.80	188.70
12		10850	4096	4.44E+07		240.24	211.38
13		11238	4096	4.60E+07		237.98	217.20
14		11625	4096	4.76E+07		255.16	211.70
15		11717	4096	4.80E+07		258.71	212.24

Block Write Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored	Disconnect	
					Enabled	Disabled
0		12400	2048	2.54E+07	143.94	142.26
1		13175	2048	2.70E+07	126.82	147.03
2		13950	2048	2.86E+07	121.64	170.63
3		14725	2048	3.02E+07	123.54	172.85
4		15500	2048	3.17E+07	124.75	175.88
5		16275	2048	3.33E+07	125.39	179.48
6		17050	2048	3.49E+07	127.28	184.55
7		17825	2048	3.65E+07	129.97	188.25
8		18600	2048	3.81E+07	130.70	190.15
9		19375	2048	3.97E+07	132.99	192.34
10		20150	2048	4.13E+07	134.97	195.08
11		20925	2048	4.29E+07	137.09	193.78
12		21700	2048	4.44E+07	137.92	199.63
13		22475	2048	4.60E+07	137.36	199.46
14		23250	2048	4.76E+07	139.06	198.64
15		23433	2048	4.80E+07	141.49	206.43

Read Cache Test 1:

Conditions	3 Pass	2 Pass	Enabled	Disabled
Disconnect			X	X
Read Cache			X	

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Disconnect		
						Enabled		Disabled
0		194	131072	2.54E+07		765.00		778.19
1		206	131072	2.70E+07		669.41		697.94
2		218	131072	2.86E+07		767.44		786.25
3		231	131072	3.03E+07		805.89		769.00
4		243	131072	3.19E+07		827.89		800.82
5		255	131072	3.34E+07		831.17		915.82
6		267	131072	3.50E+07		924.43		944.09
7		279	131072	3.66E+07		964.67		980.83
8		291	131072	3.81E+07		1016.87		1061.20
9		303	131072	3.97E+07		1007.38		1027.94
10		315	131072	4.13E+07		1000.25		1001.49
11		327	131072	4.29E+07		1017.40		1017.40
12		340	131072	4.46E+07		1098.99		1269.54
13		352	131072	4.61E+07		1092.27		1187.25
14		364	131072	4.77E+07		1150.99		1147.87
15		367	131072	4.81E+07		1134.41		1192.89

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		388	65536	2.54E+07		702.06		617.71
1		412	65536	2.70E+07		659.53		750.16
2		436	65536	2.86E+07		680.92		796.35
3		461	65536	3.02E+07		751.12		806.56
4		485	65536	3.18E+07		712.74		745.62
5		509	65536	3.34E+07		725.85		835.28
6		533	65536	3.49E+07		782.03		893.45
7		558	65536	3.66E+07		865.96		904.33
8		582	65536	3.81E+07		911.60		893.45
9		606	65536	3.97E+07		912.14		792.48
10		630	65536	4.13E+07		937.67		907.29
11		654	65536	4.29E+07		940.80		939.74
12		679	65536	4.45E+07		927.56		988.99
13		703	65536	4.61E+07		1005.18		1050.23
14		727	65536	4.76E+07		1024.39		977.07
15		733	65536	4.80E+07		1051.86		943.10

Read Cache Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		775	32768	2.54E+07		581.20		603.55
1		824	32768	2.70E+07		603.11		633.39
2		872	32768	2.86E+07		558.30		594.21
3		921	32768	3.02E+07		573.27		635.72
4		969	32768	3.18E+07		646.00		544.96
5		1018	32768	3.34E+07		692.81		682.51
6		1066	32768	3.49E+07		684.70		675.89
7		1115	32768	3.65E+07		688.14		603.72
8		1163	32768	3.81E+07		714.04		747.91
9		1211	32768	3.97E+07		731.86		766.91
10		1260	32768	4.13E+07		774.34		656.04
11		1308	32768	4.29E+07		751.45		905.97
12		1357	32768	4.45E+07		793.71		815.02
13		1405	32768	4.60E+07		767.89		662.35
14		1454	32768	4.76E+07		762.50		949.74
15		1465	32768	4.80E+07		800.27		756.31

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		1550	16384	2.54E+07		430.41		334.95
1		1647	16384	2.70E+07		425.38		370.79
2		1744	16384	2.86E+07		501.51		476.18
3		1841	16384	3.02E+07		459.53		475.02
4		1938	16384	3.18E+07		482.99		381.97
5		2035	16384	3.33E+07		491.10		500.69
6		2132	16384	3.49E+07		474.11		505.74
7		2229	16384	3.65E+07		417.86		518.22
8		2325	16384	3.81E+07		561.17		573.01
9		2422	16384	3.97E+07		555.11		475.78
10		2519	16384	4.13E+07		564.48		495.87
11		2616	16384	4.29E+07		618.07		523.33
12		2713	16384	4.44E+07		572.66		517.19
13		2810	16384	4.60E+07		547.89		602.76
14		2907	16384	4.76E+07		585.20		549.85
15		2930	16384	4.80E+07		556.61		540.77

Read Cache Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		3100	8192	2.54E+07		273.16		244.46
1		3294	8192	2.70E+07		268.46		219.58
2		3488	8192	2.86E+07		271.52		218.32
3		3682	8192	3.02E+07		285.73		209.25
4		3875	8192	3.17E+07		247.23		202.73
5		4096	8192	3.36E+07		287.41		329.24
6		4263	8192	3.49E+07		323.38		269.02
7		4457	8192	3.65E+07		302.66		334.11
8		4650	8192	3.81E+07		321.44		286.02
9		4844	8192	3.97E+07		321.43		395.67
10		5038	8192	4.13E+07		321.84		317.38
11		5232	8192	4.29E+07		343.59		325.37
12		5425	8192	4.44E+07		363.97		313.79
13		5619	8192	4.60E+07		304.82		365.05
14		5813	8192	4.76E+07		330.99		350.74
15		5859	8192	4.80E+07		344.72		374.19

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		6200	4096	2.54E+07		182.80		191.42
1		6588	4096	2.70E+07		180.64		171.41
2		6975	4096	2.86E+07		175.34		164.81
3		7363	4096	3.02E+07		161.03		173.98
4		7750	4096	3.17E+07		155.19		155.19
5		8138	4096	3.33E+07		165.92		168.66
6		8525	4096	3.49E+07		173.08		220.47
7		8913	4096	3.65E+07		171.22		223.36
8		9300	4096	3.81E+07		177.53		148.88
9		9688	4096	3.97E+07		185.04		184.65
10		10075	4096	4.13E+07		164.22		133.72
11		10463	4096	4.29E+07		167.39		211.14
12		10850	4096	4.44E+07		167.41		140.20
13		11238	4096	4.60E+07		166.08		174.17
14		11625	4096	4.76E+07		169.02		142.12
15		11717	4096	4.80E+07		160.03		197.60


Read Cache Test 1 (Continued)

Block Transfer Rate

Band Number		Number of CDBs	Block Size	Total Stored		Transfer Rate (Cache)		
						Enabled		Disabled
0		12400	2048	2.54E+07		101.79		102.20
1		13175	2048	2.70E+07		78.72		77.27
2		13950	2048	2.86E+07		92.64		56.95
3		14725	2048	3.02E+07		81.47		97.84
4		15500	2048	3.17E+07		85.50		80.23
5		16275	2048	3.33E+07		85.64		88.66
6		17050	2048	3.49E+07		86.50		103.16
7		17825	2048	3.65E+07		90.69		82.42
8		18600	2048	3.81E+07		92.64		111.32
9		19375	2048	3.97E+07		92.85		113.42
10		20150	2048	4.13E+07		95.44		82.98
11		20925	2048	4.29E+07		96.17		86.77
12		21700	2048	4.44E+07		97.60		91.97
13		22475	2048	4.60E+07		97.58		119.58
14		23250	2048	4.76E+07		93.45		126.75
15		23433	2048	4.80E+07		88.95		86.04


Pinnacle Sierra Average Write Rate (2 Pass)

■ Disconnect Enabled ◊ Disconnect Disabled




Pinnacle Sierra 2-Pass Write Cache. 128k Block

■ Disconnect Disabled ■ Disconnect Enabled




Pinnacle Sierra 2-Pass Write Cache, 64k Block

■ Disconnect Disabled ■ Disconnect Enabled




Pinnacle Sierra 2-Pass Write Cache, 32k Block

■ Disconnect Disabled ■ Disconnect Enabled




Pinnacle Sierra 2-Pass Write Cache, 16k Block

□ Disconnect Disabled ■ Disconnect Enabled




Pinnacle Sierra 2-Pass Write Cache, 8k Block


■ Disconnect Disabled □ Disconnect Enabled




Pinnacle Sierra 2-Pass Write Cache, 4k Block




Pinnacle Sierra 2-Pass Write Cache, 2k Block




Pinnacle Sierra 2-Pass Band Write (Disconnect Disabled)




Pinnacle Sierra 2-Pass Band Write (Disconnect Enabled)




Pinnacle Sierra Disconnect Effects, 128k Block Write




Pinnacle Sierra Disconnect Effects, 64k Block Write




Pinnacle Sierra Disconnect Effects, 32k Block Write




Pinnacle Sierra Disconnect Effects, 16k Block Write





Pinnacle Sierra Disconnect Effects, 8k Block Write




Pinnacle Sierra Disconnect Effects, 4k Block Write






Pinnacle Sierra Disconnect Effects, 128k Block Read




Pinnacle Sierra Disconnect Effects, 64k Block Read




Pinnacle Sierra Disconnect Effects, 32k Block Read

—■— Disabled —◇— Enabled




Pinnacle Sierra Disconnect Effects, 16k Block Read


—■— Disabled —◇— Enabled



Pinnacle Sierra Disconnect Effects, 8k Block Read




Pinnacle Sierra Disconnect Effects, 4k Block Read



Pinnacle Sierra Disconnect Effects, 2k Block Read

—■— Disabled —◆— Enabled



**MISSION
OF
ROME LABORATORY**

Mission. The mission of Rome Laboratory is to advance the science and technologies of command, control, communications and intelligence and to transition them into systems to meet customer needs. To achieve this, Rome Lab:

- a. Conducts vigorous research, development and test programs in all applicable technologies;
- b. Transitions technology to current and future systems to improve operational capability, readiness, and supportability;
- c. Provides a full range of technical support to Air Force Material Command product centers and other Air Force organizations;
- d. Promotes transfer of technology to the private sector;
- e. Maintains leading edge technological expertise in the areas of surveillance, communications, command and control, intelligence, reliability science, electro-magnetic technology, photonics, signal processing, and computational science.

The thrust areas of technical competence include: Surveillance, Communications, Command and Control, Intelligence, Signal Processing, Computer Science and Technology, Electromagnetic Technology, Photonics and Reliability Sciences.